AI-based consumption forecast to reduce energy costs for the operation of charging infrastructure in retail

Kolja Eger, Nick Krüger
Faculty of Engineering and Computer Science
Hamburg University of Applied Sciences (HAW Hamburg), Germany
{kolja.eger, nick.krueger}@haw-hamburg.de

Nils Heinrich
Envidatec GmbH
Hamburg, Germany
nils.heinrich@envidatec.com

***Abstract*—The buildup of the charging infrastructure in retail significantly changes the load profiles of these energy consumers resulting in higher costs due to power peaks. This paper proposes a new approach for energy management at supermarkets where the cooling processes are used as flexibility. The approach makes use of the time gaps between charging processes to selectively intensify the cooling processes. This energy reserve is used when new charging processes begin. Key capability is a forecast module based on deep learning. The proposed CNN-LSTM model with additional input signals for seasonality and public holidays shows good performance for a short-term prediction over two hours.**

***Index Terms*—charging infrastructure, energy consumption, supermarkets, consumption forecasts, neural networks, CNN-LSTM**

I. Introduction

Companies in industry and commerce are electrifying their vehicle fleets and building publicly accessible charging infrastructures for customers to attract more and to create new business opportunities. Retail chains are leading the way, currently equipping their stores with fast charging solutions up to 400 kW. This increases the peak power demand of a single store easily by a factor of two or more resulting in higher costs for electricity, especially for grid charges which depend on the peak power demand. Additionally, required grid expansions result in extra costs or could not be realized by the grid operators (in a timely manner).

The EcoCharge project therefore pursues a novel approach that predicts the energy consumption in retail stores with high accuracy in order to coordinate the processes in the store efficiently with the power requirements of the charging infrastructure. This is achieved by a novel energy management system based on AI technologies. Thereby, throttling of the charging process is avoided, as this leads to lower revenues as well as reduced customer acceptance. Furthermore, conventional methods like peak shaving are not suitable, since situational switch off and time shifts of specific loads can disrupt operational processes and worsen overall energy efficiency. In particular, our approach makes use of the thermal storage of refrigeration systems, which is often already available in large quantities in the retail sector. This avoids additional investments, e.g. in electrical storage systems like batteries.

The proposed solution addresses supermarkets with charging stations. Supermarkets are volatile energy consumers whose energy demand is strongly based on the refrigeration processes. The energy efficiency depends heavily on the operating status and can vary rapidly over a wide range depending on the load situation. The proposed approach makes use of the time gaps between charging processes to selectively intensify the cooling processes in the market in advance of an emerging peak demand. This builds up an energy reserve which can be used during peak demand without impacting grid charges. Thereby, two objectives are met: Firstly, load peaks are not only cut off as with common approaches like peak shaving but are also managed without any negative effects on operational processes. Secondly, unfavorable operating conditions are avoided so that the efficiency of the systems can be stabilized at a permanently high level.

For the proposed peak load optimization a short-term consumption forecast is required to efficiently build up the cooling reserve when free charging capacities are identified and before new charging processes begin. With high utilization of the charging infrastructure (e.g. 70 to 80%), only short time intervals remain, so that the peak load must be known at least two hours in advance. The energy demand of supermarkets can vary by 20 to 30% at short notice during the course of the day. Since neither the processes in the market nor their interactions are fully known, the consumption profile cannot be modeled analytically. Thus, the proposed forecast is based on a novel deep learning algorithm, which interacts with a process control system that accesses the market's refrigeration systems and controls their performance by specifying target values (see Fig. 1).

The focus of the paper is on the AI-based forecast model. In Section II we briefly summarize related work and in Section III the used data sets are described and analyzed. The neural network model for the AI-based approach is presented in Section IV. Section V shows its parametrization and the results of the paper, followed by a conclusion in Section VI.

This work is part of the research project "EcoCharge" funded by Hamburgische Investitions- und Förderbank (IFB Hamburg).

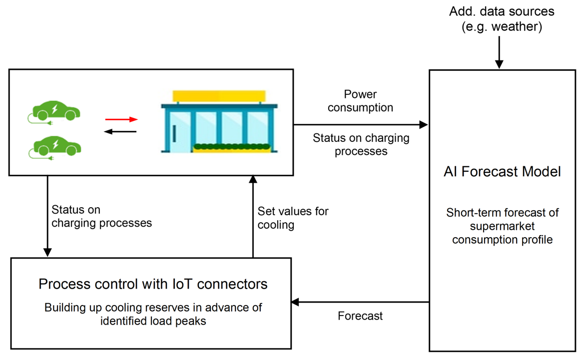

Fig. 1. Schematic model for AI-based Energy Management System

II. Related Work

The energy management of supermarkets is typically based solely on statistical data evaluation in an EMS (energy management system). Energy-saving functions are provided for the operational optimization of refrigeration systems, which are limited to timer programs and simple logical operations between process variables. The savings potential is relatively low at approx. 10%.

Approaches based on AI are discussed in research but are not widely used in the market yet. We are not aware of any publication on the combined approach to optimize refrigeration system and charging infrastructure together. Only with a focus on one of the subsystems data-driven approaches based on neural networks are proposed.

In [1] four different methods (e.g. neural networks and support vector machines) are used to investigate the charging behavior of wallboxes on the low-voltage grid. In [2] the required charging infrastructure for Park&Ride is evaluated using Random Forest methods. In [3] simple neural networks are used to predict the charging behavior (average charging rate and power profile) of electric cars. Further research on machine learning (ML) for the planning of charging infrastructure has been compiled in a literature review [4].

The energy consumption of supermarkets is also predicted and optimized with the help of machine learning in various studies. In [5] four regression models (including neural networks and support vector machines) are trained on historical consumption data from supermarkets in order to predict the consumption of new stores with certain characteristics. [6] investigates the control of refrigeration systems using Q-learning, an approach from reinforcement learning.

According to the literature review in [7] support vector machines (SVM), Gaussian-based regression and clustering algorithms are the most commonly used ML-based methods to predict energy consumption for buildings in general.

For energy consumption and time-series forecasting in general, Recurrent Neural Networks (RNNs) are often the method of choice. [8] covers multiple different RNN-based forecasting methods using Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU). LSTM can also be combined with Convolutional Neural Networks (CNN) resulting in an architecture with one-dimensional convolutional layers in combination with LSTM layers. These CNN-LSTM models show promising results with lower error rates than other neural networks on data sets such as gas field production [9], stock prices [10] and gold prices [11]. CNN-LSTM is also used by the authors in [12] for day-ahead forecasting of electricity consumption in Germany.

III. Energy Data for Supermarkets

A. Consumption data

In this work we use data from two real supermarkets in Northern Germany. The two supermarkets differ in type and consumption. The year of construction is in the 1990s (with modernizations in 2021) and in 2009, respectively. The consumption data is available in a quarter-hourly resolution for two and three years, respectively. The data set is complete and contains only valid values ("no NaNs").

The average yearly energy consumption is 1,200 MWh and 1,650 MWh with peak power demand of 300 kW and 344 kW, respectively.

Data on electricity consumption shows complex seasonality with daily, weekly, monthly and also annual patterns as well as days with unusual consumption (often on and around public holidays) [13]. This is also the case for the supermarket data where we observe recurring patterns over the course of the day, and between weekdays, Saturdays and Sundays.

Figure 2 shows heatmaps for different seasonal periods for one of the supermarkets. Figure 2a shows the average power value for a quarter of an hour over the course of the day and day of the week. Power demand is low on Sundays as well as at night. It increases before opening hours (i.e. 7 am) with peaks in the early morning and decrease after the market is closed (i.e. 9 pm).

In Figure 2b the average power per day over the years 2021 until mid of 2024 is depicted. This plot shows unusual times of consumption at the end of 2021 and in the first months of 2022 which is a major difference of this supermarket compared to the other data set. The reasons for this unusual consumption period is not fully clarified. In addition, lower consumption can be seen on public holidays (e.g. on Christmas or Eastern) as well as higher consumption on specific Sundays (e.g. 10.9. and 5.11.2023).

For the other supermarket the power consumption on different time scales is depicted as box plots in Figure 3. Also here, daily and weekly seasonality can be clearly seen. A monthly seasonality is not strongly pronounced. One possible reason for this could be the additional energy consumption for air conditioning in these supermarkets during the summer months, in addition to the typically high consumption in winter.

B. Simulated charging data

The available data sets do not include the power consumption for the charging processes at the charging stations of the supermarkets. Therefore, a simulation model is used to

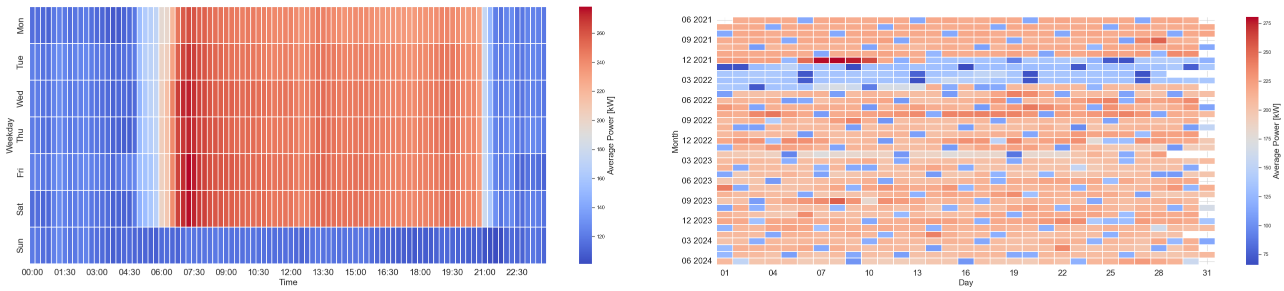

(a) Average power over day and weekdays

(b) Average power per day

Fig. 2. Power consumption data for Supermarket 1

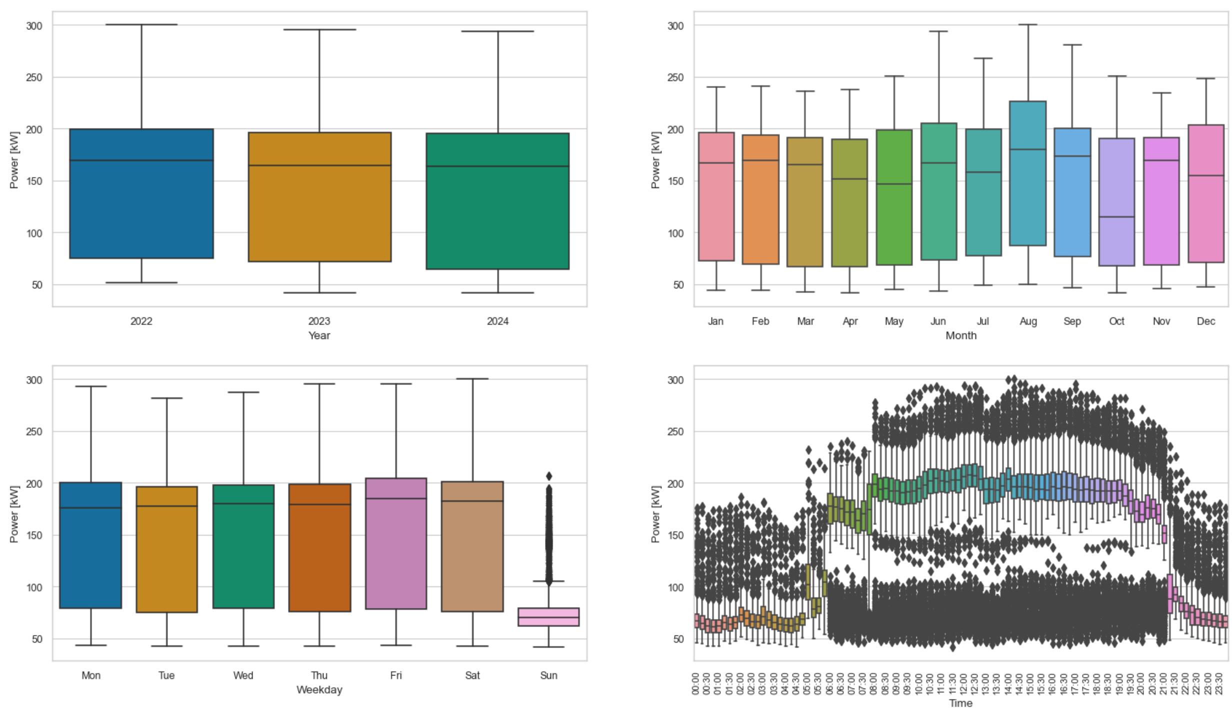

Fig. 3. Power consumption data of Supermarket 2

create load profiles for the aggregated charging processes. The general consumption profiles of the supermarkets discussed in Section III-A are augmented with this simulated data to train the model as described in Section IV.

This procedure is seen as an intermediate step until real data on the charging processes are available over a longer period of time. Furthermore, it ensures early validation of the prediction approach presented.

The simulation model is described in detail in [14] and follows a similar approach as in [15], [16]. It includes parameters on the supermarket itself, the charging stations and the electric vehicles (EVs). For supermarkets, the opening times are taken into account, as not all charging stations are accessible outside opening hours, for example. The number of charging points and the maximum charging capacity are parameters on the charging stations. The charging processes are described by the arrival process, the charging duration and the initial state-of-charge (SoC) of the EV. For the arrival process and the charging duration different distributions can be modeled, e.g. Poisson and Weibull based e.g. on typical values for parking durations in retail [17]. The simulation model includes also charging profiles of different types of EVs which are chosen randomly during simulation. Currently, the simulation model neglects the effect of the temperature during the charging process. As an example the simulated load profile over one day is depicted in Figure 4. It shows the number of charging processes and the load profile during the opening hours for one day. The load profile shows short peaks of different magnitude which are caused mainly by the typical charging profile of the different types of EVs.

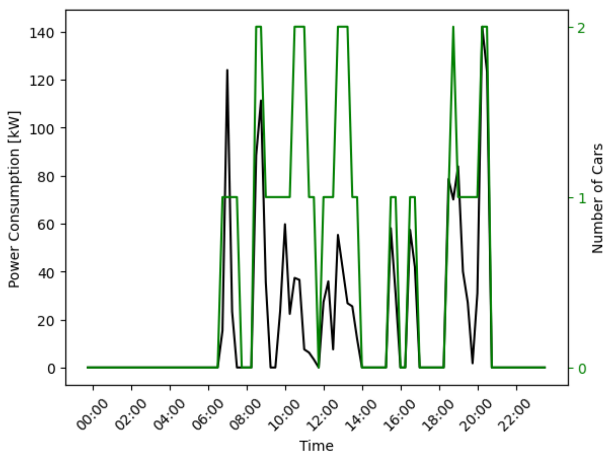

Fig. 4. Load profile for simulated charging processes

IV. Neural Network Models

As discussed in the related work section, the CNN-LSTM network architecture has shown promising results at forecasting time series data, including creating day-ahead forecasts of the energy consumption of Germany, as discussed in [12]. In order to create accurate power consumption forecasts, we use this architecture as well as additional input signals to model seasonal trends.

First, the input data for each forecast is fed into the input layer of the network. In our case, this data consists of the power consumption of the supermarket for a number of previous input time steps. We also create additional input features to help the model understand the seasonal effects of the data set. These additional features consist of sine and cosine values for each time step, with a period of an hour, a day, a week, a month, and a year, as well as a signal indicating the day of the week, and public holidays. The holiday signal uses a public calender of German holidays and is set to 1 on public holidays, 0.5 on the days before a public holiday and 0 on all other days. Any combination of these input features can be fed into the network for any number of input time steps, both being parameters for the network optimization. A possible combination of these input features is shown in Figure 5, using the day curves as well as the day of the week and holiday signals. In that figure, the time steps in the green range are part of the input window, which in the shown case contains 96 values and 5 features. The values in the red range are part of the output window, which is what the model is trained to forecast for the given input window.

The hidden layers following the input layer of the network are one-dimensional convolutional layers. In these layers, a convolutional kernel slides across the features of the previous layer along the temporal axis, resulting in a feature map. The weights of the kernel are adjusted by the model during training. For each convolutional layer, the optimizable hyperparameters are the number of convolutional filters per layer, the size of the kernel and the step size. In this model

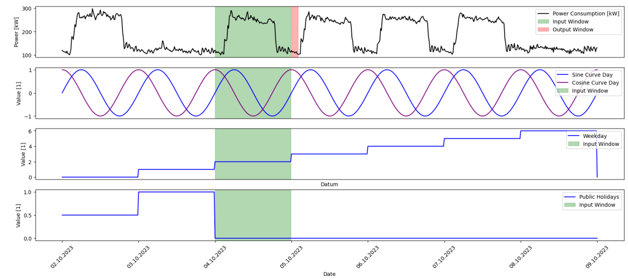

Fig. 5. Example of input and output windowing for 96 input time steps

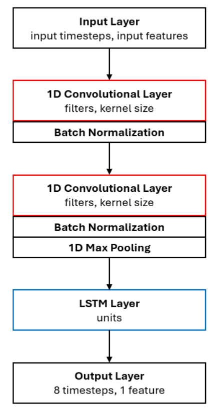

Fig. 6. CNN-LSTM model architecture with optimizable parameters

implementation, each convolutional layer is followed by a batch normalization layer. After the convolutional layers, the resulting feature maps are fed into a maximum pooling layer and then a long short-term memory (LSTM) layer. This is a recurrent layer capable of learning short and long-term temporal trends via each cell of a layer passing the resulting cell state and hidden state for each time step along to the next time step. The hyperparameters for this layer type are the number of such LSTM cells per layer. The outputs of the final LSTM layer are then fed into a fully-connected output layer with a linear output function. For this task, as we want to forecast 2 hour time frames and our data is sampled every 15 minutes, the output layer contains 8 units. Figure 6 showcases the resulting CNN-LSTM architecture.

V. Parameter Optimization & Performance

In order to explore the effects of different input and network parameters on the forecasting accuracy and to find the overall best performing model, we implemented, trained, and evaluated various models of this architecture using a grid search algorithm, where we tested every possible combination of input and network parameters and measure their performance.

Table I shows the parameters for this search, as well as the number of models resulting from these combinations. These models are individually trained and evaluated using both a validation data set and a testing data set. The error rates used are the mean absolute error (MAE), the mean absolute percentage error (MAPE), and the mean squared error (MSE).

A. Parameter Optimization Results

In order to optimize the network parameters, we ran two searches using these search parameters. Both searches use the data of supermarket 1, with one also using the simulated power consumption of two charging stations. The training data ranges from 02.06.2021 until 30.06.2023, the validation data from 01.07.2023 until 31.12.2023 and the testing data from 01.01.2024 until 30.06.2024. The two plots in Figure 7 show the forecasting performance of all trained models in these two searches.

Based on the search results, we selected a network architecture to forecast the supermarket and charging power consumption. The parameters of the models were chosen by analyzing the effect of each parameter on the network performance across the search, selecting the best possible combination. The resulting best parameters are the same for both searches. Table II contains the resulting parameters of the selected network.

After parameter selection, we trained the selected architecture 20 times for supermarket 1 and supermarket 2, both with and without adding simulated charging data. This was done in order to confirm if our results are reproducible, as the random initialization of the network weights can cause a difference in forecasting performance when training a model with the same hyperparameters and input data multiple times.

As the data set of supermarket 2 is smaller, the training data set for this supermarket covers time steps from 29.07.2022 until 30.06.2023. The validation and testing data sets cover the same time frames as for supermarket 1 in order to create comparable results. Table III contains the mean values and the 95% confidence intervals for the different error measures.

The resulting confidence intervals for the MAPE of both models are within one percent point. Therefore, our models perform reliable and can be re-trained without a significant loss of accuracy. The models of supermarket 1 outperform the models of supermarket 2, both with and without charging data. As the time frames of the validation and testing data sets are the same across both supermarkets, this can be explained by the smaller amount of training data used to train the models of supermarket 2. Finally, for both supermarkets, the error rates of the model increase with the addition of simulated charging data. Evidently, the randomness of the vehicle arrival process leads to a loss of forecasting performance. To improve the performance of our networks in this area will be a major focus of our future work.

B. Seasonality

As our approach utilizes additional input features to model seasonality, this is a particular area of interest when evaluating the forecasting performance of the networks. For this reason, we create heatmaps to showcase the mean performance of the networks across months, days and time steps. Figure 8 shows the mean error rate of two models trained on data from supermarket 1, both with and without simulated charging data, across the days on the validation and testing data sets. Additionally, Figure 9 shows the mean error rate of the same two models across the day of the week and time of day. In both cases, the first few days of the validation and testing data set are not visualized, as for these days, there are not enough time steps in the two data sets respectively to create a forecast.

The resulting heatmaps show multiple interesting trends, such as the large error rates of both models during the week of Christmas in 2023 and Easter 2024. The addition public holidays signal as an input therefore appears to not have the same increase in forecasting performance as shown on other data sets [12]. This may be explained by the smaller training data sets, which leads to the model encountering fewer public holidays during training. In addition, the models also show an increased error rate during Sundays with exceptionally high power consumption, such as the 05.11.2023. This effect, as well as a slightly worse performance during Sundays in general, leads to Sundays being the weekday with the worst forecasting performance. Overall, the model with additional simulated charging data performs worse than the model forecasting the data from only the the supermarket. In addition, it performs worse during opening than closing hours on weekdays, whereas the model without additional charging data performs worse during closing hours. This effect is due to the model struggling to forecast the random nature of the charging processes, which only takes place during opening hours. The performance of the two models during closing hours on weekdays, where no charging processes happen, is comparable. This loss in forecasting accuracy due to the charging processes is amplified on Sundays, yielding larger error rates than on weekdays, as the power consumption of the supermarket is lower.

VI. Conclusion

This work proposes a novel approach to manage peak power demand for supermarkets with charging infrastructure. We use the cooling processes in the supermarket as flexibility to alleviate the high load peaks of the charging processes. A short-term load forecast is required to build up cooling reserves before the charging of newly arriving EVs begins. The forecast is realized with a deep learning model based on CNN-LSTM with additional input signals for seasonality and public holidays.

The forecast model achieves very good results for supermarkets without charging processes, esp. with increasing amount of training data. The MAPE is below 6% with training data of near to two years. The error rate increases when less data is available for training. This can be seen for a second data set with an error rate of 9% with training data of less than one year. In addition, the prediction accuracy decreases if the data for charging processes are taken into account (11 and 16%, repsectively). For this, we see different reasons: On the one hand, charging processes are highly stochastic with steep increase at the beginning of charging processes. The arrival process and the duration are mainly driven by the customers'

TABLE I
CNN-LSTM SEARCH PARAMETERS

Parameter	Values
Input Time Steps	96, 288
Fixed Input Columns	Power Consumption [kW]
Dynamic Input Columns	Day, Week, Month, Years Curves, Day of the Week, Holidays
Number of CNN Layers	2
CNN Filters	16, 32
CNN Kernel Size	2, 4
Number of LSTM Layers	1
LSTM Units	128, 256
Total Number of Models	8192

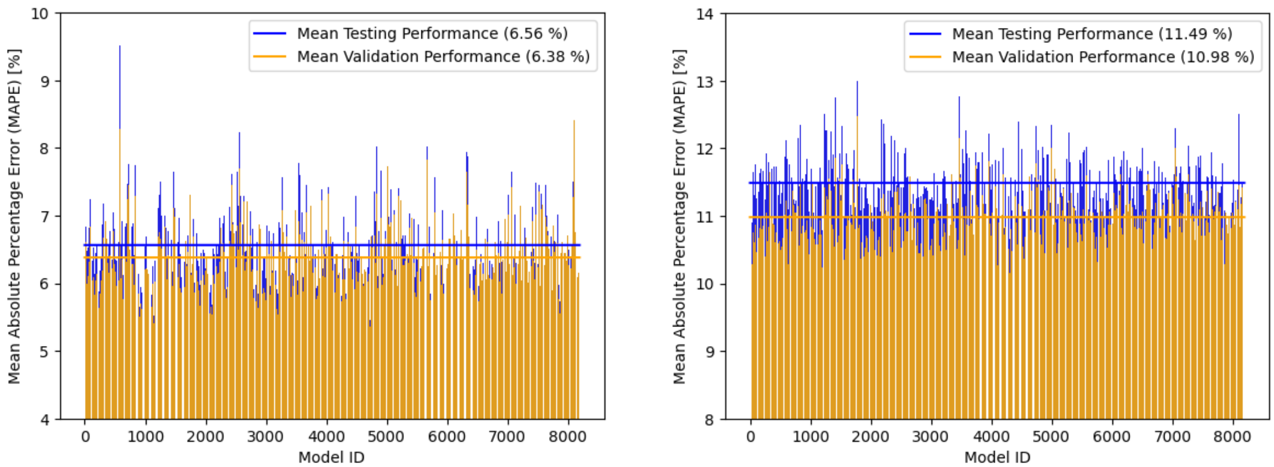

(a) Search without simulated charging data

(b) Model search with simulated charging data

Fig. 7. Parameter optimization results for Supermarket 1

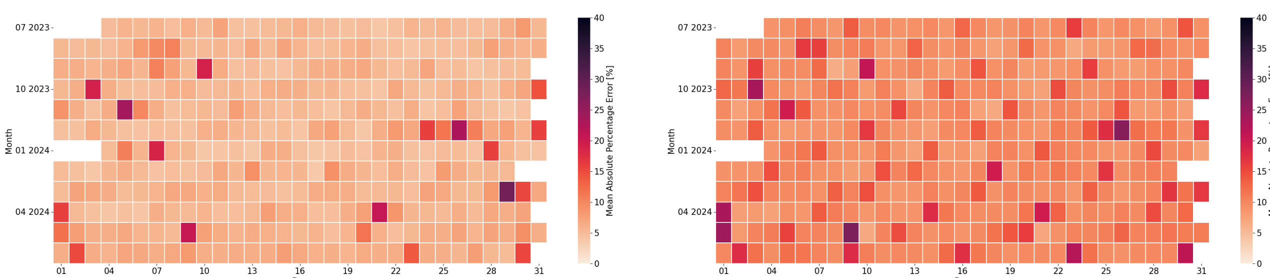

(a) Without charging data

(b) With charging data

Fig. 8. Average forecasting MAPE per day for Supermarket 1

usage behavior. Typically charging profiles depend on the type of the car and vary with changing SoC. On the other hand, simplifications were made in the used simulation model. For example, the rate of the arrival process is constant during opening hours of the supermarket and does not include any variations or patterns over the day or week.

In future work, we like to replace the simulation model for the charging processes with real data. This data is currently being collected and a data history sufficient for the first experiments is expected this year. In addition, our forecast model could be evaluated for additional supermarkets. Data collection and provision for this is also under discussion with associated partners in the EcoCharge project. Based on the promising results presented in this work we plan to integrate the forecast model with the process control as depicted in Figure 1.

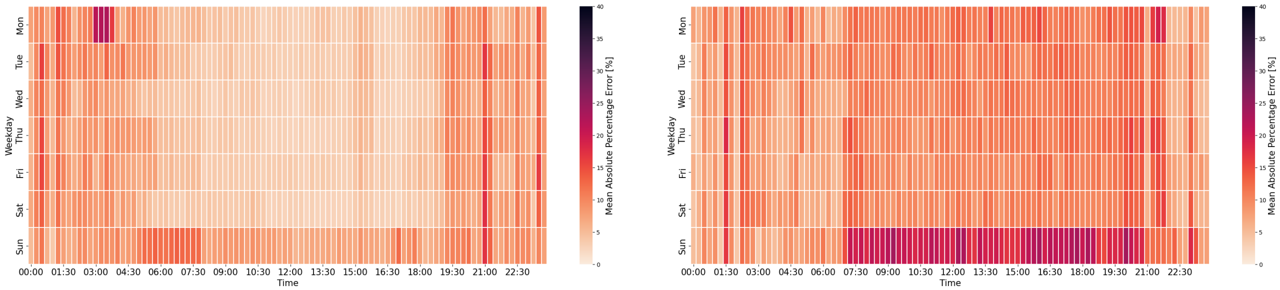

(a) Without charging data

(b) With charging data

Fig. 9. Average forecasting MAPE over time and weekdays for Supermarket 1

TABLE II
SELECTED CNN-LSTM PARAMETERS

Parameter	Best values
Input Time Steps	288
Dynamic Input Columns	Day, Day of the Week, Holidays
Number of CNN Layers	2
CNN Filters	(16, 16)
CNN Kernel Size	(4, 2)
Number of LSTM Layers	1
LSTM Units	256

TABLE III
NETWORK CONFIDENCE INTERVALS FOR THE VALIDATION DATA SET

	Supermarket 1		Supermarket 2	
	-	charging	-	charging
Upper 95% CI MAE [kW]	9.86	23.57	10.28	23.14
Mean MAE [kW]	10.06	23.69	10.47	23.32
Lower 95% CI MAE [kW]	10.25	23.82	10.65	23.49
Upper 95% CI MAPE [%]	5.64	10.43	9.09	15.40
Mean MAPE [%]	5.77	10.52	9.29	15.61
Lower 95% CI MAPE [%]	5.90	10.61	9.48	15.83
Upper 95% CI MSE [kW]	14.40	33.58	14.61	32.80
Mean MSE [kW]	14.66	33.66	14.89	32.98
Lower 95% CI MSE [kW]	14.91	33.75	15.16	33.16

REFERENCES

[1] P. Eitel and P. Stolle, "A machine learning approach to model the future distribution of e-mobility and its impact on the power grid," *Energy Inform*, vol. 5 (Suppl 1), 31, 2022. [Online]. Available: https://doi.org/10.1186/s42162-022-00203-w

[2] M. Ferrara, C. Liberto, M. Nigro, M. Trojani, and G. Valenti, "Multimodal choice model for e-mobility scenarios," *Transportation Research Procedia*, vol. 37, pp. 409–416, 2019, 21st EURO Working Group on Transportation Meeting, EWGT 2018, 17th – 19th September 2018, Braunschweig, Germany. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2352146518306264

[3] A. Ramachandran, A. Balakrishna, P. Kundzicz, and A. Neti, "Predicting electric vehicle charging station usage: Using machine learning to estimate individual station statistics from physical configurations of charging station networks," 2018. [Online]. Available: https://arxiv.org/abs/1804.00714

[4] S. Deb, "Machine learning for solving charging infrastructure planning problems: A comprehensive review," *Energies*, vol. 14, no. 23, 2021. [Online]. Available: https://www.mdpi.com/1996-1073/14/23/7833

[5] R. Granell, C. J. Axon, M. Kolokotroni, and D. C. Wallom, "Predicting electricity demand profiles of new supermarkets using machine learning," *Energy and Buildings*, vol. 234, p. 110635, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0378778820334216

[6] A. Beghi, M. Rampazzo, and S. Zorzi, "Reinforcement learning control of transcritical carbon dioxide supermarket refrigeration systems," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 13 754–13 759, 2017, 20th IFAC World Congress. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2405896317334924

[7] S. Seyedzadeh, F. Rahimian, I. Glesk, and M. Roper, "Machine learning for estimation of building energy consumption and performance: a review," *Vis. in Eng.*, vol. 6, 2018.

[8] H. Hewamalage, C. Bergmeir, and K. Bandara, "Recurrent neural networks for time series forecasting: Current status and future directions," *International Journal of Forecasting*, vol. 37, no. 1, pp. 388–427, 2021.

[9] W. Zha, Y. Liu, Y. Wan, R. Luo, D. Li, S. Yang, and Y. Xu, "Forecasting monthly gas field production based on the cnn-lstm model," *Energy*, vol. 260, p. 124889, 2022.

[10] W. Lu, J. Li, Y. Li, A. Sun, and J. Wang, "A cnn-lstm-based model to forecast stock prices," *Complexity*, vol. 2020, pp. 1–10, 2020.

[11] I. E. Livieris, E. Pintelas, and P. Pintelas, "A cnn–lstm model for gold price time-series forecasting," *Neural computing and applications*, vol. 32, pp. 17 351–17 360, 2020.

[12] N. Krüger, K. Eger, and W. Renz, "Smardcast: Day-ahead forecasting of german electricity consumption with deep learning," in *2024 International Conference on Smart Energy Systems and Technologies (SEST)*. IEEE, 2024, pp. 1–6.

[13] J. W. Taylor, "Triple seasonal methods for short-term electricity demand forecasting," *European Journal of Operational Research*, vol. 204, no. 1, pp. 139–152, 2010. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S037722170900705X

[14] J. Kindel, "Entwicklung eines Simulationsmodells mit Python für das Ladeverhalten von Elektromobilität," Bachelor thesis, HAW Hamburg, 2024.

[15] T. A. Hertlein, T. Blenk, C. Weindl, and J. Ochs, "Object-oriented charging model for the simulation of grid-serving intelligent charging infrastructure," in *ETG Congress 2023*, 2023, pp. 1–8.

[16] M. Gilleran, E. Bonnema, J. Woods, P. Mishra, I. Doebber, C. Hunter, M. Mitchell, and M. Mann, "Impact of electric vehicle charging on the power demand of retail buildings," *Advances in Applied Energy*, vol. 4, p. 100062, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2666792421000548

[17] EHI Retail Institute e.V. Elektromobilität im Handel 2024. Accessed on Jan. 8th, 2025. [Online]. Available: https://www.ehi.org/produkt/whitepaper-elektromobilitaet-im-handel-2024/