38th International Electric Vehicle Symposium and Exhibition (EVS38)
Göteborg, Sweden, June 15-18, 2025

Integration of Nowcasting in Electric Vehicle Charging Infrastructures

Sascha Gohlke[1], Zoltan Nochta[2]
[1]*Karlsruhe University of Applied Sciences (corresponding author), sascha.gohlke@h-ka.de*
[2]*Karlsruhe University of Applied Sciences, zoltan.nochta@h-ka.de*

Executive Summary

To operate electric vehicle (EV) fleets in a safe and efficient manner, many companies have been deploying charging infrastructures (CI) at their premises. A CI is typically monitored and controlled by a software system that can dynamically steer, schedule, and prioritize charging processes to satisfy EV charging demand. Forecasting of different system parameters of a CI, such as how many charging points will be occupied during the day, can help create accurate charge plans. In this paper, we deal with the applicability of continuous Nowcasting, i.e., frequently executed short-term forecasts, to predict the next few data points based on past and current situation in a CI. In our experiments, we forecast the number of connected EVs over a rolling two-hour horizon, employing weighting schemes to emphasize recent observations and multi-step strategies to forecast each prediction step. We then demonstrate how these predictions can be integrated to improve an example CI.
Keywords: charging business models, consumer demand, energy management, smart-charging, smart grid integration and grid management

1 Introduction

According to [1], renewable energy sources accounted for over 95% of net capacity additions in 2023 worldwide, incl. 72% solar and 20% wind energy. These energy sources are characterized by high volatility, which poses a significant challenge in properly balancing energy grids. In order to have sufficient electrical energy available at any time, excess energy can be stored and consumed at times of lower generation, or energy consumption can be shifted to times of higher generation. Companies often have high energy demand, but also the flexibility to shift electricity consumption in a controlled manner, depending on their business processes. The usage of electric vehicles (EVs), incl. cars, delivery vans, forklifts, trucks, buses, heavy duty machinery, etc., in the business context helps increase this flexibility in general. The International Energy Agency estimated that 99% of all high-performance EV chargers will be operated at company sites with a total installed capacity of 2,000 GW by 2035 [2]. A company that runs one or more charging infrastructures (CIs) at its locations, such as production sites, warehouses, logistic hubs, distribution centers, office buildings, etc., can implement flexibility measures through the dynamic (periodic or event-triggered) determination and assignment of consumable power to each CI. Note, that the consumable power for EV charging can be set to a lower value than the maximum technically possible power limit in the specific CI. The task of a local EV charging management system is to distribute the allocated consumable power among ongoing charging sessions.
A company can optimize the cost of energy for its EV fleet by estimating the respective energy demand day-ahead and purchasing energy on the market accordingly. In addition, it can also flexibly manage (limit) the actual power consumption in each CI during the day, depending on the actual demand and intra-day energy prices as well [3]. This strategy not only requires the planning of EV charging processes

for every business day upfront, but also the continuous monitoring of actual demand and the capability to change power assignments if needed.

In context of CI, a charging plan that was created for a given day is expected to be followed until significant changes occur. For example, due to unexpected bad weather, the power demand in the morning hours can be much higher than estimated the day before, because more employees would get to work by car (EV) as usual. Besides such unexpected events, the limited availability of proper data, e.g., when an e-truck that needs charging will arrive at a warehouse location, can also lead to inaccurate day-ahaed forecasts. To mitigate such deviations in a timely manner, continuous monitoring of the demanded power is needed. However, monitoring itself can only detect issues that have already occurred: In the above example the charging management system would realize in the morning that much more EVs are connected than expected and try to distribute the available (limited) power among them. A few minutes later the system would realize that the number of connected EVs has even grown, and it would, again, create a new schedule for (all) charging sessions in order to keep the power limit foreseen for the morning hours.

In this paper, we propose a Nowcasting approach to continuously observe and frequently, e.g., every 15 minutes, predict charging point (CP) utilization throughout the day. The short-term predictions can help detect emerging trends in the CI early, for example, that the number of connected EVs starts to deviate from previously expected values. In reaction, the available power for EV charging can dynamically be adapted (increased or decreased) and cover the the demand of all newly arriving vehicles.

The remainder of the paper is organized as follows: In Section 2 a definition of Nowcasting and a brief overview of relevant related work is given, followed by the summary of the research methodology in Section 3. Section 4 describes requirements and our technical approach to periodically predict the number of used charging stations within a CI. Section 5 is dedicated to the experiments and the prediction results based on a real- world dataset. In Section 6 we discuss how prediction results can actually influence the energy management related decisions. Finally, in Section 7 we provide a conclusion and outline next steps.

2 Related Work

Nowcasting originated from meteorology and has advanced from primitive temporal extrapolation of meteorological radar and satellite imagery to advanced systems and algorithms that process data from a variety of sources [4, 5]. In economics, Nowcasting refers to the process of forming expectations about key features of the economy, monitoring and incorporating new data releases in real time and revise estimations whenever the realizations diverge from the initial predictions [6, 7].

We adapt the basic idea of Nowcasting and define it in the context of company EV charging as follows: Nowcasting in company EV charging infrastructure refers to the integration of the current (”now”) system state with continuously updated short-term forecasts of key variables, such as charge point utilization, energy demand, or EV arrival and departure times. It provides a real-time view of anticipated near-future conditions, so-called future system states, to support timely operational decision making. In other words, nowcasting refers to the process of describing what the system will look like shortly based on its current state and behavior. [8].

To ensure smooth operation in a CI, smart charging systems (SCS) are developed to manage the power distribution between charging stations and ensure the safety of the CI by balancing loads to prevent fuses and transformers from being overloaded [9, 10]. The overall efficiency of such measures is highly dependent on the data available to the SCS. The required information is often stored in external systems, such as energy management systems, vehicle back-ends, etc., and must be made available. In addition to protecting the CI, the SCS can also contribute to cost reduction by redistributing power consumption to periods when energy prices are low by establishing an upper software limit for the maximum power available for charging at times of high energy prices. A power limit is usually applied in combination with a smart charging algorithm, which distributes power to charging sessions by prioritizing sessions based on various factors, taking into account EV and CI hardware limitations [11, 12].

3 Methodology

This work addresses the problem of adapting to unusual trends in a CI that are difficult to anticipate during day-ahead planning. By forecasting the number of connected EVs for a horizon up to 2 hours at a resolution of 15 minutes, we allow decision making to make timely adjustments to the power limitation of the CI. We implement and compare two distinct predictive models: Long Short-Term Memory (LSTM), a recurrent neural networks suited for sequence learning and XGBoost, a tree-based ensemble regressor trained via gradient boosting. To generate the short-term forecasts, we use two combinations of data sets. Each combination contains a two and a four calendar week training set, followed by one validation and one testing week. The use of calendar weeks enables the capture of cyclic temporal patterns in workweek-related behavior. All datasets are derived from a real-world dataset collected from a smart charging system used to manage the CI of a company [9]. The models are enriched with temporal features and optimized through hyperparameter tuning using Optuna. In addition, we apply three different weighting strategies (linear, squared, and exponential) to assign increasing importance to the most recent observations of the current day. Additionally, we propose three multi-horizon forecasting strategies (Direct, DirRec, and MIMO) to generate predictions over a two-hour forecast horizon. The experimental evaluation assesses the effectiveness of various combinations of weighting schemes and multi-horizon forecasting strategies based on the Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE).

4 Nowcasting Concept in the Context of EV Charging

We consider the following requirements important for the calculation and integration of short-term forecasts in EV charging systems. Note that numbering is only provided for referencing purposes and does not imply any order of importance.

1. **Suitable forecasting interval and horizon:** Predictions about the future status of the CI are enablers of efficient management procedures during the day, incl. planning of charging sessions, limiting charging power for charging stations, etc. We consider a forecasting horizon of two hours in advance with a prediction recalculation frequency of 15 minutes as suitable. The forecasting horizon of two hours is selected as a compromise between prediction accuracy, as forecast reliability tends to degrade over time, and a sufficient (re-)planning period for related decision making. We consider two hours as sufficient to take measures against a detected, unusual trend, e.g., a higher number of connected EVs than is typical at the given time of the day. The selection of a 15-minute forecasting interval is aligned to the typical interval for updating energy prices in intra-day energy markets.

2. **Real time integration and computation time**: Short-term predictions require the gathering of relevant data in near real-time, fast computation using prediction models and a fast subsequent decision making process as well. In the case of a CI, the time from detecting a trigger event to the roll-out of an updated charge plan assigned to the charging stations should be minimized.

3. **Adaption to dynamic usage patterns:** Prioritizing the most recent data enables to adapt to usage patterns that deviate from regular behavior, e.g., as a result from changing weather conditions, special events, holidays, traffic congestion, etc. Assigning higher weights to more recent data has the effect of biasing the prediction models towards recent data. This decreases the influence of outdated patterns, and predictions can adapt to short-term trends more accurately. However, an overly aggressive weighting strategy can result in the exclusion of significant long-term patterns. In the case of the CI belonging to a company with regular daytime working hours, it may be beneficial to weight data points that belong to the specific working day. If the CI belongs to a company with continuous operations, for example a 24/7 logistics hub, day-specific patterns may be of less significance. In this case, it may be more beneficial to set a fixed number of most recent data points to be weighted.

4. **Fast model adaption and deployment**: Using a limited amount of input data for prediction leads to reduced training and deployment times as well as fast model updates for real-time adaption. It allows companies to quickly deploy and iterate on forecasting solutions without the need for extensive historical data collection. In addition, using only a few weeks of historical data further helps capture most recent patterns.

4.1 Approach to increase the significance of most recent data points

Short-term forecasts are used to predict n data points in the near future, based on periodically captured past data points of an existing time series dataset, as shown in Figure 1.

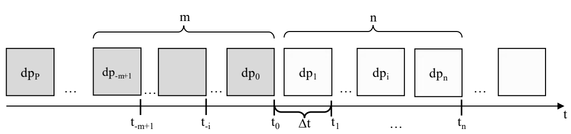

Figure 1: Prediction of future data points in a time series based on past and current observations

Note that a data point can represent a single system parameter, but also a collection of different measurement values at time t_i, such as the number of ongoing charging sessions, the power drawn, the total capacity of charging EVs' batteries, etc. The equidistant data points dp_i with $P \leq i \leq 0$ contain data about P past system states captured at t_i respectively, whereas the data point dp_0 represents the latest (current) known state at t_0. The data points dp_i with $0 < i \leq n$ are future data points to be predicted according to the selected frequency. In a practical setup, for example, with $\Delta t = 5min$, $P = 120$ and

$n = 4$, the 120 latest data points (captured in a ten hours time-frame) would be used to predict the next four data points occurring within the next twenty minutes.
In order to increase the significance of the most recent data points for the predictions, w_i weights ($-m + 1 \leq i \leq 0$) are computed and assigned to m most recent data points ($m \leq P$), such that dp_i will be weighted with the corresponding w_i value.
The applied weighting ensures that newer data points have a greater influence on the prediction of the next few future data points compared to older data points. The underlying basic assumption is that a CI behaves as an inertial system with regard to parameters, such as the number of ongoing charging sessions: The duration of charging sessions is relatively long compared to Δt, and it is likely that only a small number of EVs arrive or depart within the next few observation/capturing periods.
We consider the following weighting strategies to determine and assign weights to m selected past data points m:
The linear (L) approach, illustrated by one possible implementation in Equation (1), assigns weights that linearly increase with for more recent observations, see also Table 1.

$$w_i = m + i \tag{1}$$

The squared (S) approach, illustrated by one possible implementation in Equation (2), quadratically increases the weights of data points, leading to a higher weight for the most recent observations compared to the linear approach.

$$w_i = (m + i)^2 \tag{2}$$

The exponential (E) approach, illustrated by one possible implementation in Equation (3), increases the weight of data points in much more aggressively.

$$w_i = e^{\frac{q}{m}(m+i)} \tag{3}$$

The growth rate $\frac{q}{m+i}$ is set based on the number of data points to weight by adapting the parameter q. Parameter m represents the number of data points selected for weighting. This ensures that that numerical overflow is avoided when weighting a full day. At the same time, it ensures a higher weight for the most recent data point than the other approaches even for small numbers of m, as shown in Table 1.

Table 1: Example weights for $m = 8$ data points calculated by different approaches

	$i = 0$	-1	-2	-3	-4	-5	-6	-7
$L(w_i)$	8	7	6	5	4	3	2	1
$S(w_i)$	64	49	36	25	16	9	4	1
$E(w_i)$	22026.47	6310.69	1808.04	518.01	148.41	42.52	12.18	3.49

Models that allow the application of weights to single data points using the $sample_weight$ parameter, weights can be incorporated to adjust model training accordingly. In addition, weights can be incorporated by using the Weighted Moving Average (WMA) as an additional input feature. If required, the weights can be dynamically adjusted based on the performance of the prediction model [13, 14].

$$\text{WMA}_{i_0} = \frac{\sum_{i=-m+1}^{0} w_i \cdot dp_{i_0+i}}{\sum_{i=-m+1}^{0} w_i} \tag{4}$$

WMA_{i_0} denotes the WMA at the current time index i_0, based on the most recent m observations. It is computed using the values x_{i_0+i} for $i \in [-m + 1, 0]$, where each point is assigned a weight w_i based on the selected approach (see L, S, E above) for weighting. The weighted sum is then normalized by the sum of the weights. For example, if $m = 8$, the weighted moving average at the current time index i_0 is computed using the most recent eight data points: $dp_{i_0-7}, dp_{i_0-6}, \ldots, dp_{i_0}$.
The combination of the proposed exponential weighting approach with the WMA results in an exponential WMA that is conceptually related to the time series prediction algorithm ARIMA(0,1,1) [15].

4.2 Multi-Step Strategies for Prediction Model Training

To predict a selected number of future equidistant data points with a given frequency (e.g., $\Delta t = 5min$), we examine different multi-step strategies [16, 17]. A so-called single-step strategy computes a specific future data point by indirectly considering all past data points that were used to train the prediction model, as shown in Figure 2. The prediction model itself remains static, i.e., there is no retraining applied.

In multi-step forecasting strategies (see in Figure 2) multiple subsequent future data points are predicted. Thereby, the model can use previously predicted data points as input data (together with known past values) to predict further future data points. As an alternative, predicted values can be added to the training dataset, and the model can be retrained accordingly, before executing the next prediction.

- The Direct strategy trains an independent model for each prediction horizon by using already known past observations as training dataset. The approach helps minimize prediction errors for the target value at the specific horizon. The main disadvantage is the increasing training effort with increasing number of prediction steps.
- The Recursive strategy trains a single model that uses previous predictions as inputs during inference to generate multi-step forecasts.
- The Direct-Recursive (DirRec) strategy combines the advantages of both the Direct and the Recursive strategies, by training independent models for each forecast horizon, as in the Direct strategy, and by incorporating previous predictions as inputs during inference as in the Recursive strategy [18]. For short-term forecasting of unusual, volatile changes in CP utilization, we do not use the Recursive strategy as it is prone to error accumulation and may fail to adapt to evolving patterns in the predicted values.
- To mitigate stepwise error accumulation, the Multi-Input-Multi-Output (MIMO) strategy predicts a single output vector containing all prediction steps simultaneously based on the latest observations. Another advantage of this strategy is that only a single model needs to be trained.

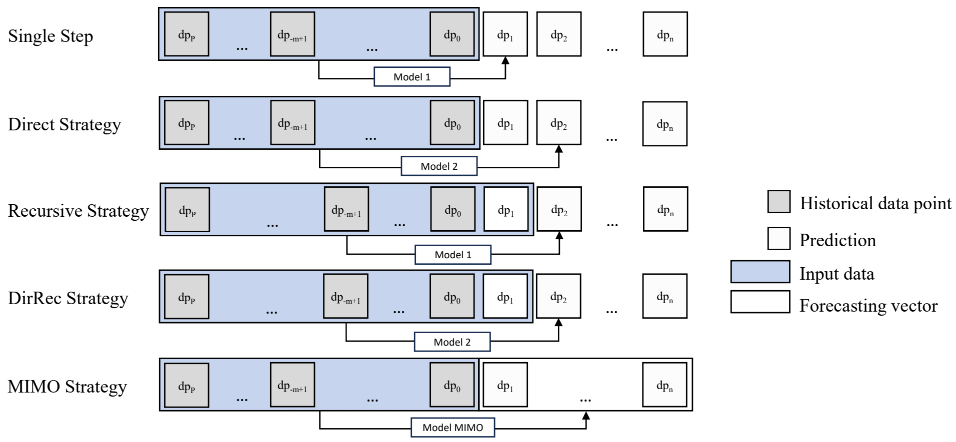

Figure 2: Overview of forecasting strategies for multiple future time steps.

5 Nowcasting Experiments to Predict Active Charging Sessions

To detect short-term changes, we predict the number of connected EVs and focus mainly on three of the requirements defined in Section 4. To address requirement 1, we apply the multi-horizon forecasting strategies Direct, DirRec and MIMO in our experiments. A fixed horizon of two hours, corresponding to eight predicted equidistant data points is used to validate the strategies. Furthermore, we address requirement 3, by applying the proposed linear, squared and exponential weighting schemes to increase importance of most recent data points (see Section 4.1) and adapt to dynamic usage patterns. The selected weighting scheme is based on data points of the current day, meaning that the number of data points to weight increases with every prediction step until the end of the day. To address requirement 4, we evaluate fast model adaption and deployment by evaluating the impact of input data limited to data points from two and four weeks.

5.1 Data Set Selection

Data used for experiments shows characteristic workweek patterns from EV charging behavior related to typical working hours. Most EVs arrive in the morning and depart in the afternoon with additional trends, such as volatility in utilization around lunchtime. In addition, typical weekday-related behavior is observable: On Wednesday there are usually less EVs compared to other working days, and there are only very few (sometimes zero) charging sessions on the weekend. However, there are also other things to consider, sometimes EVs are connected to the charging station over night, sometimes even for multiple days, see Figure 4 from September 20 to 22. We did not exclude such sessions from the dataset as they can for instance, be used for V2G applications and the detection of such behavior can be beneficial. To utilize the workweek behavior in the data, the data sets are split in training, testing and validation data of full work weeks. To evaluate the impact of the duration of the training set, Datasets combinations are created with periods of two and four weeks sharing a common test and validation week.

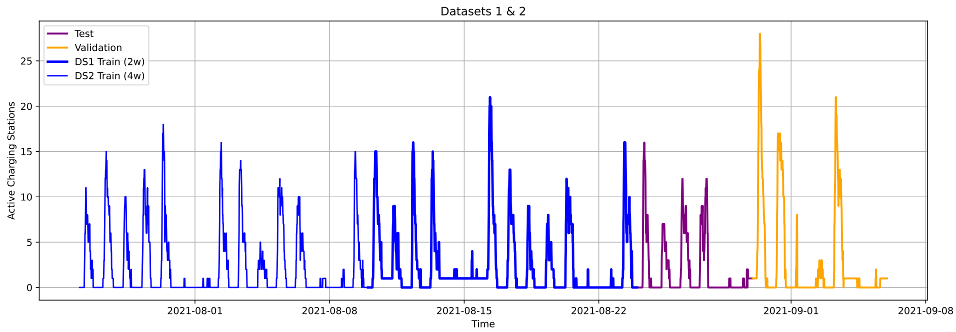

Figure 3: Two training data sets spanning two (DS1) and four regular business weeks (DS2). The validation data represents an irregular week.

Figure 3 shows the two datasets DS1 and DS2 that share a common test (illustrated by the purple curve) and validation week (illustrated by the yellow curve) and. DS1 contains data points from two weeks for training as illustrated by the bold blue curve. DS2 contains four weeks of training data, including data points from DS1, as illustrated by the thinner blue curve.

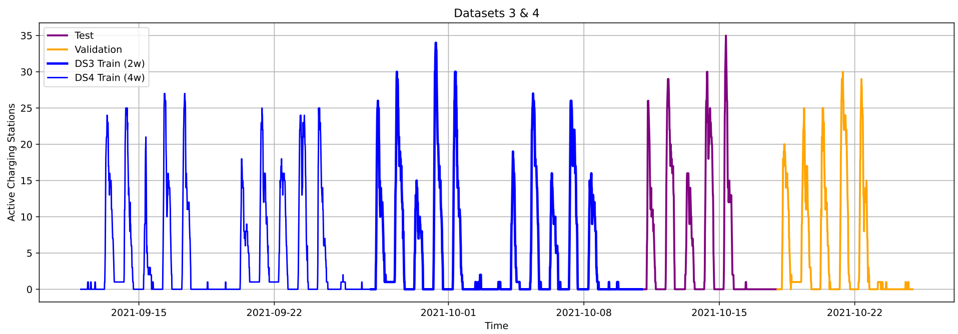

Figure 4: The two sets of training data, 2 weeks (DS3) and 4 weeks (DS4), showing regular business activity and one irregular week as validation data

Figure 4 shows DS3 and DS4 using the same color schema for training, test and validation data as in Figure 3. Similarly, DS3 consists of a two-week training period that is included in the four-weeks training period of DS4, with both data sets sharing the same test week.

5.2 Model Selection and Optimization

We use two different methods to generate short-term predictions in our experiments. First, Extreme Gradient Boosting (XGBoost) is used as it offers strong performance for the training of multiple models as required in multi-horizon strategies such as the Direct and DirRec strategy. It is further extensively used and recommended in Literature to capture short-term changes [19, 20]. Second, LSTM is selected because of its design to capture temporal dependencies and sequential patterns and its common use in predictions related to forecasting in the are of EV charging, as shown in an overview of different forecasting methods [21]. Feature selection and hyperparameter tuning were optimized and validated using a one-week test set and a rolling window approach with a single step forecasting strategy. The only source for additional feature extraction was the timestamp associated with each data point. For the XGBoost model we use features that include lagged values (`lag_1`, `lag_4`, `lag_96`, `lag_672`), time-based cyclical encodings (`dayofweek_sin`, `dayofweek_cos`, `hour_sin`, `hour_cos`), and the binary indicator `is_weekend`. Features for LSTM are normalized and the cyclical nature of temporal features is addressed by a sine and cosine transformation. It includes the target variable `target`, time-based

cyclical encodings (minutes_sin, minutes_cos, dayofweek_sin, dayofweek_cos), the flag is_weekend, and the weekly lag lag_96. All Features are scaled to $[-1, 1]$.
Both models are trained using the Adam optimizer with mean squared error (MSE) as the loss function. Hyperparameter optimization is performed with Optuna and results are shown in Table 2. Parameters are optimized for each combination of XGBoost and LSTM with each dataset. Due to the different $output\_length$, using eight instead of one for the MIMO strategy, additional parameters are used. Note that for LSTM, we set $epochs = 1$ based on findings in [22].

Table 2: Best hyperparameter for LSTM and XGBoost

Model	**DS**	Input Length	Units/ Estimators	Learning Rate	Batch Size	Dropout/ Gamma	Max Depth
LSTM	1	4	32	0.01	32	0	–
LSTM	2	2	128	0.01	96	0	–
LSTM	3	2	64	0.01	32	0.2	–
LSTM	4	2	16	0.01	96	0	–
LSTM (MIMO)	1	32	128	0.01	32	0.2	–
LSTM (MIMO)	2	32	8	0.01	16	0.2	–
LSTM (MIMO)	3	2	8	0.01	32	0.2	–
LSTM (MIMO)	4	4	16	0.01	96	0.2	–
XGBoost	1	–	100	0.105	–	0.2	3
XGBoost	2	–	120	0.248	–	0.048	8
XGBoost	3	–	600	0.312	–	0.178	4
XGBoost	4	–	1310	0.006	–	0.123	6

Following the optimization phase, the validation data set is utilized by the rolling window to apply the multi-step strategies. For XGBoost, the proposed time-dependent weighting schemes were applied by computing a weight for each training instance based on the data point of the current day. These weights were then integrated into the training process via the $sample\_weight$ parameter to increase importance of most recent patterns. While this introduces additional computational overhead, the short input sequences and the efficiency of XGBoost make this approach feasible for applying dynamic weighting schemes. However, XGBoost is not included in the evaluation of the MIMO strategy, because it is designed for single-target regression. In contrast, LSTM requires more training time. To address this, we included our weighting strategies by adding the WMA as a feature, as described in 4.1. This allows training of the horizon-specific LSTM models once at the start of the test phase, instead of every timestep, to reduce computation time.

5.3 Experimental Results

As expected, experimental results indicate a decline in prediction accuracy with increasing forecast horizon. Figure 5 shows the predictions of the XGBoost model for a full week, generated using linear weighting and the Direct multi-horizon strategy applied to the data set DS2.
To enhance visibility of single prediction steps, only step 1 (15 minutes), 4 (1 hour) and 8 (2 hours) are displayed. An example of the adaption to trends in most recent data is Monday, August 30, a day with significant higher utilization as usual. This case illustrates the impact of short-term forecasting under atypical conditions, showing the influence of past patterns in the more distant steps. The more recent a future data point is, the better the model adapts to the ongoing trend. Predictions of more distant future data points tend to revert toward the typical utilization patterns in the training data. A similar pattern is observed in the next day, in which the utilization remains high for longer than usual. On Thursday, September 2, the pattern is reversed, with significantly fewer charging sessions than usual. Initial forecasts reflect typical usage patterns, but as new data points are included, predictions are adapted to the new pattern by predicting lower utilization of CPs.
Figure 6 illustrates different prediction steps on Monday August 30, generated by the LSTM model using linear weighting and the Direct multi-horizon strategy applied to DS2. As noted earlier, utilization on that day exceeded typical levels. Predictions capture the higher utilization trend, estimating higher than usual utilization for the day.
Table 3 and 4 show the performance of the LSTM and XGBoost models across all data sets, using the Direct strategy and with Exponential (E), Linear (L), and Square (S) weighting schemes. RMSE and MAE for steps between Table 3 for step 1 and Table 4 for step 8 generally increase for each step. Results show that the first forecasting step of XGBoost trained on DS1 with two weeks of data result in

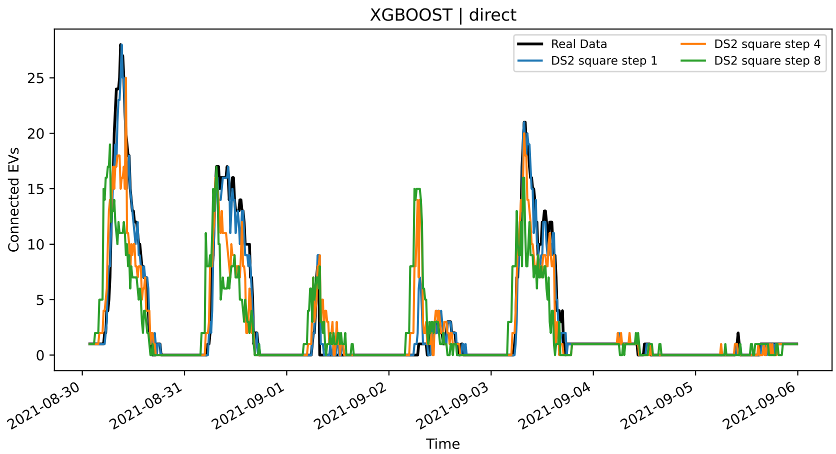

Figure 5: Results based on DS2 utilizing the Direct multi-horizon strategy with linear weighting predicted by the XGBoost model.

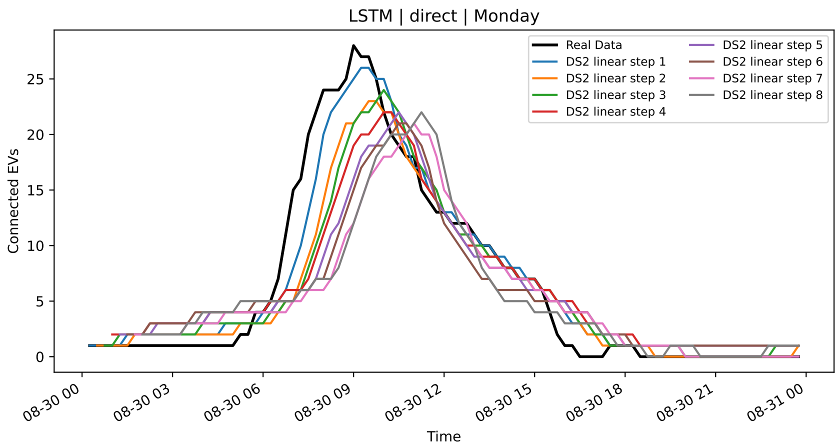

Figure 6: Results based on DS2 utilizing the Direct multi-horizon strategy with linear weighting predicted by the LSTM model.

	XGB(E)	LSTM(E)	XGB(L)	LSTM(L)	XGB(S)	LSTM(S)
(DS1, RMSE)	1.022	1.268	0.906	1.261	0.910	1.203
(DS2, RMSE)	1.110	1.114	1.061	1.080	1.053	1.103
(DS3, RMSE)	1.510	1.393	1.496	1.482	1.442	1.340
(DS4, RMSE)	1.207	1.097	1.253	0.957	1.216	1.288
(DS1, MAE)	0.445	0.905	0.355	0.866	0.383	0.798
(DS2, MAE)	0.490	0.547	0.448	0.624	0.444	0.672
(DS3, MAE)	0.706	0.878	0.698	0.853	0.691	0.761
(DS4, MAE)	0.578	0.674	0.599	0.586	0.577	0.667

Table 3: Comparison of the performance of LSTM and XGBoost at the first step (15 minutes) for the Direct strategy across datasets DS1–DS4 using RMSE and MAE.

slightly better RMSE than forecasts trained on DS2 with four weeks of data. In contrast, for the data sets containing regular business weeks, the opposite is the case. The four week training set DS4 has better RMSE values for all weighting schemes than DS3. For LSTM, the opposite trend is observed, for both combinations of the data sets.

	XGB(E)	LSTM(E)	XGB(L)	LSTM(L)	XGB(S)	LSTM(S)
(DS1, RMSE)	3.011	3.023	2.813	3.111	2.926	3.097
(DS2, RMSE)	3.217	3.130	3.106	3.150	3.197	3.177
(DS3, RMSE)	3.142	4.469	3.132	4.311	3.152	4.357
(DS4, RMSE)	2.454	4.678	2.502	4.560	2.469	4.314
(DS1, MAE)	1.552	1.659	1.337	1.785	1.423	1.737
(DS2, MAE)	1.498	1.350	1.414	1.353	1.475	1.538
(DS3, MAE)	1.580	2.350	1.556	2.240	1.562	2.266
(DS4, MAE)	1.237	2.461	1.258	2.408	1.230	2.332

Table 4: Comparison of the performance of LSTM and XGBoost at the last step (2 hours) for the Direct strategy across datasets DS1–DS4 using RMSE and MAE.

	XGB(E)	LSTM(E)	XGB(L)	LSTM(L)	XGB(S)	LSTM(S)
(DS1, RMSE)	3.280	3.082	2.943	3.260	3.037	3.200
(DS2, RMSE)	3.046	3.015	3.120	3.054	3.089	2.959
(DS3, RMSE)	25.516	4.361	3.827	4.327	4.063	4.436
(DS4, RMSE)	2.664	4.489	2.608	4.841	2.720	4.556
(DS1, MAE)	1.550	1.925	1.394	1.938	1.403	1.961
(DS2, MAE)	1.383	1.344	1.418	1.256	1.389	1.341
(DS3, MAE)	2.842	2.336	1.739	2.356	1.820	2.386
(DS4, MAE)	1.330	2.303	1.332	2.260	1.351	2.380

Table 5: Comparison of the performance of LSTM and XGBoost at the last step (2 hours) for the DirRec strategy across datasets DS1–DS4 using RMSE and MAE.

Table 5 shows the performance of the LSTM and XGBoost models across all data sets, using the DirRec strategy. A special measure is the high RMSE of the XGBoost model in DS3 with exponential weighting. The value is the result of accumulated errors in the last horizon on Thursday morning, where each overestimation adds exponentially to the next, leading to predictions that are significantly too high. Overall results indicate that unusual usage patterns can be detected, especially at shorter forecast horizons. The best performance based on RMSE results from the XGBoost model utilizing the direct strategy with linear weighting of most recent observations.

6 Integration and Utilization of Predictions in Charging Infrastructures

To show the impact of the Nowcasting approach, we take an example scenario shown in Figure 7.
It is assumed that an energy supplier will announce energy prices day-ahead, with prices being higher in the morning and lower in the afternoon. The blue dotted curve illustrates the power consumption that would result from unrestricted EV charging. The introduced power limit, shown by the black dashed curve, is aligned with the expected dynamic energy prices. Limiting the power leads to unsatisfied energy demand, as illustrated by the blue-dotted area above the limitation. The missing energy is shifted to be charged by the EVs during the afternoon by increasing the power limit as soon as energy prices drop. This results in a higher power demand up to the increased power limit, until the required amount of energy has been charged by the EVs (see the blue-dotted area below the power limitation).
However, such an approach lacks the ability to react to changes that occur during execution and may result in inadequate charging of EVs in cases with deviation from the expected daily pattern. For example, a customer event can lead to a higher number of EVs, and, thus, to higher power consumption than expected, as illustrated in Figure 8.
The blue dotted curve in Figure 8 illustrates power consumption that would occur in absence of power restrictions on a day with high CP utilization. The power demand on this day is slightly higher as usual, and the decrease at lunchtime is less pronounced. The red curve shows the power needed to compensate the previously undelivered energy resulting from earlier power limitations. When EVs depart in the afternoon, less energy as required is delivered, so that EVs leave with insufficiently charged batteries. To minimize energy that could not successfully be shifted (see the white area between the red and the blue

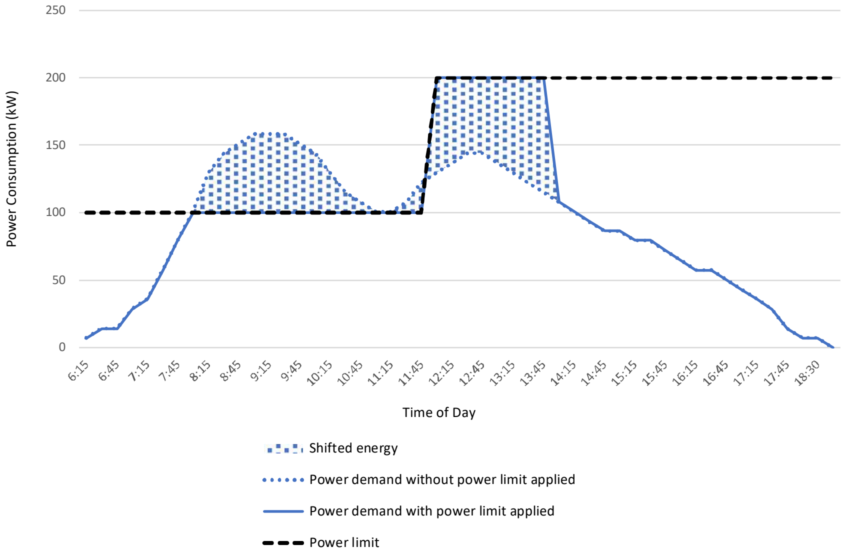

Figure 7: Power consumption of a company's CI, impacted by power limits to reduce energy costs

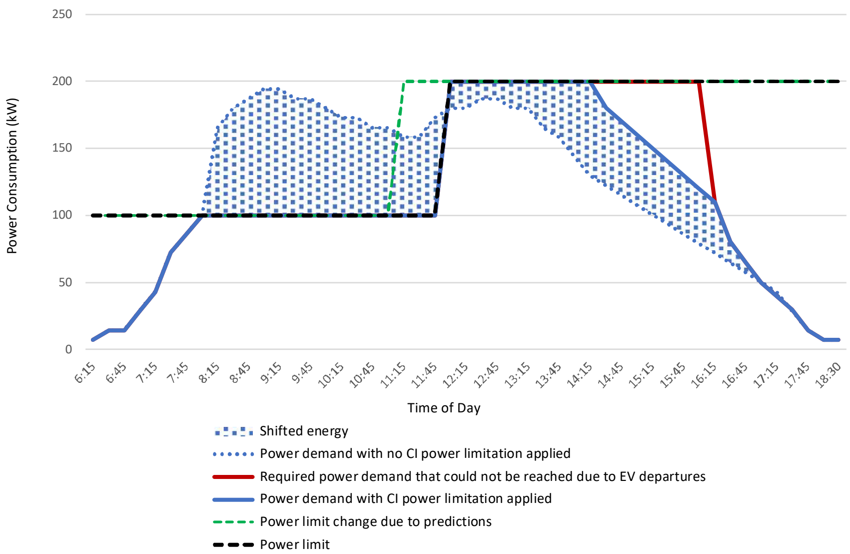

Figure 8: Power consumption of a company's CI on a day with higher CP utilization than expected.Based on nowcasting results, the power limit can be increased earlier as planned.

curve), it is essential to identify deviations from the planned charging schedule at an early stage. This is accomplished by running short-term predictions. For each predicted future data point we calculate the deviation between the estimated day-ahead CP utilization and the number of predicted EVs at the current time (now). The deviations observed in the predicted data points (in our case $n = 8$ in a two-hour forecast horizon) indicate that the actual number of connected EVs has been diverging from the estimation. This leads to a new trend in CP utilization, that can be included in decision making for further power planning, for example, to re-adjust the maximum available and thus distributable power for charging. We assume that adjusting the CI power limit, constitutes a necessary precautionary measure, as deviations from the planned charging power can lead to the emergence of more costly scenarios. In the case of a lower than expected trend, the company needs to distribute excess energy bought the day ahead energy by increasing the power consumption of other local assets, otherwise it can pay fines to its energy provider . If the identified trend is higher than expected as it is the case in our scenario, the company can decide if EV charging under the new conditions is still sufficient or if an adjustment to the power limit is required. Without appropriate adjustments, situations may arise in which insufficiently charged EVs require an additional stop for recharging during their next trip, potentially leading to missed delivery deadlines or preventing employees from attending important customer appointments. To prevent such cases in the proposed scenario, Figure 8 shows the adjustment of the power limitation by increasing the power limitation early, as indicated by the green curve. As a result, the available power is now distributed more effectively among active charging sessions, enabling a greater amount of energy to be delivered throughout the day and ensuring improved SoC on EV departure. Possible measures a company can take to increase the available charging power include reducing the power consumption of other local assets or purchasing additional energy from its energy provider.

7 Conclusion and Future Work

In this work, we proposed a Nowcasting approach to adapt to unusual trends in a CI by generating short-term forecasts to predict CP utilization. To improve responsiveness to recent changes, we introduced linear, squared, and exponential weighting schemes. For multi-horizon forecasting, we employed the Direct, DirRec, and MIMO strategies, each designed to produce multiple forecasts across different future time steps. Experiments with LSTM and XGBoost, incorporating these techniques, indicate that forecasting up to 2 hours ahead at a 15-minute resolution is beneficial. Finally, we outlined how these forecasts can be integrated into company decision-making processes and discussed their potential advantages in real-world applications.
Future work includes several directions to further evaluate and enhance our approach. For instance, we plan to assess the impact of expanding the input data scope, such as by incorporating data spanning multiple months. We also aim to validate our assumption that aligning the data split with business weeks leads to better performance than conventional percentage-based splits. Additionally, we intend to investigate the effect of varying the number of data points used in the weighting schemes. Specifically, we will compare our current approach of applying weights dynamically to data points within the same day, with fixed-length windows, such as using data points from the last four or eight hours. Additionally, we intend to explore whether reducing the time interval between data points to, for example, five minutes remains computationally feasible. Furthermore, we plan to implement and adapt the approach in simulation environments to enable real-time forecasting as EVs arrive at or depart from the CI. Integrating data as soon as it becomes available, such as the detection of an arriving vehicle is crucial to further develop the Nowcasting approach.

Acknowledgements

We sincerely thank Volha Hubashyna for her help in developing the concepts and for carrying out initial experiments that laid the foundation for this work. As part of the project *BiFlex-Industrie* [1] this work has been funded by the German Federal Ministry of Economic Affairs and Climate Action (Grant number: 01MV23020A).

References

[1] A. P. Teixeira and L. Demôro, “Energy Transition Factbook 2024,” BloombergNEF, Presentation, Oct. 2024.

[2] IEA, “Global EV Outlook 2024,” International Energy Agency, Tech. Rep., Apr. 2024.

[3] K. Schert, F. Biedenbach, T. Müller, M. Kluge, and Z. Nochta, “On the Aggregation and Monetization of Flexible Loads in the Context of EV Fleets,” *World Electric Vehicle Journal*, vol. 15, no. 5, p. 216, May 2024.

[4] World Meteorological Organization (WMO), *Guidelines for Nowcasting Techniques*. Geneva: WMO, 2017, no. WMO-No. 1198.

[5] C. Mass, C. F. Mass *et al.*, “Nowcasting: The next revolution in weather prediction,” *Bulletin of the American Meteorological Society*, pp. 1–23, 2011.

[6] M. Bańbura, D. Giannone, M. Modugno, and L. Reichlin, “Now-Casting and the Real-Time Data Flow,” in *Handbook of Economic Forecasting*, ser. Handbooks in Economics, G. Elliott and A. Timmermann, Eds. Amsterdam, Netherlands: North Holland, Jan. 2013, pp. 195–237.

[7] J. Cimadomo, D. Giannone, M. Lenza, F. Monti, and A. Sokol, “Nowcasting with large Bayesian vector autoregressions,” *Journal of Econometrics*, vol. 231, no. 2, pp. 500–519, Dec. 2022.

[1] https://www.h-ka.de/en/isrg/projects/biflex-industrie

[8] S. Gohlke and Z. Nochta, "Towards a short-term forecasting framework to efficiently charge company EV fleets," in *7th E-Mobility Power System Integration Symposium (EMOB 2023)*. Copenhagen, Denmark: Institution of Engineering and Technology, Sep. 2023, pp. 191–198.

[9] T. Fleck, S. Gohlke, and Z. Nochta, "A System for the Efficient Charging of EV Fleets," *World Electric Vehicle Journal*, vol. 14, no. 12, p. 335, 2023. [Online]. Available: https://www.mdpi.com/2032-6653/14/12/335

[10] K. V. Sastry, S. Holla, S. Tater, E. Gustafson, D. G. Taylor, and M. J. Leamy, "Design and Demonstration of a Smart Charging System for Plug-in Electric Vehicles," in *2023 IEEE Transportation Electrification Conference & Expo (ITEC)*, 2023, pp. 1–6. [Online]. Available: https://ieeexplore.ieee.org/document/10186905

[11] O. Frendo, N. Gaertner, and H. Stuckenschmidt, "Open Source Algorithm for Smart Charging of Electric Vehicle Fleets," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 9, pp. 6014–6022, Sep. 2021.

[12] M. F. Voss, S. P. Haveman, and G. M. Bonnema, "In-Company Smart Charging: Development of a Simulation Model to Facilitate a Smart EV Charging System," *Energies*, vol. 14, no. 20, p. 6723, 2021. [Online]. Available: https://www.mdpi.com/1996-1073/14/20/6723

[13] X. Hu, X. Tang, Q. Zhang, Z. Chen, and Y. Zhang, "Combined machine learning forecasting method for short-term power load based on the dynamic weight adjustment," *Energy Reports*, vol. 9, pp. 866–873, Sep. 2023.

[14] M. Xiang, S. Du, J. Yu, and Y. Xu, "Short-term load forecasting based on sample weights assignment," *Energy Reports*, vol. 8, pp. 783–791, Nov. 2022.

[15] D. McCarthy and S. T. Jensen, "Power-weighted densities for time series data," *The Annals of Applied Statistics*, vol. 10, no. 1, pp. 305–334, Mar. 2016.

[16] G. Bontempi, S. Ben Taieb, and Y.-A. Le Borgne, "Machine Learning Strategies for Time Series Forecasting," in *Business Intelligence: Second European Summer School, eBISS 2012, Brussels, Belgium, July 15-21, 2012, Tutorial Lectures*, M.-A. Aufaure and E. Zimányi, Eds. Berlin, Heidelberg: Springer, 2013, pp. 62–77.

[17] S. Ben Taieb, G. Bontempi, A. F. Atiya, and A. Sorjamaa, "A review and comparison of strategies for multi-step ahead time series forecasting based on the NN5 forecasting competition," *Expert Systems with Applications*, vol. 39, no. 8, pp. 7067–7083, Jun. 2012.

[18] A. Sorjamaa and A. Lendasse, "Time series prediction using dirrec strategy," in *ESANN 2006 proceedings*, vol. 6, 01 2006, pp. 143–148.

[19] E. Aguilar Madrid and N. Antonio, "Short-Term Electricity Load Forecasting with Machine Learning," *Information*, vol. 12, no. 2, p. 50, Feb. 2021.

[20] O. Frendo, N. Gaertner, and H. Stuckenschmidt, "Improving Smart Charging Prioritization by Predicting Electric Vehicle Departure Time," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 10, pp. 6646–6653, Oct. 2021.

[21] M. Rashid, T. Elfouly, and N. Chen, "A Comprehensive Survey of Electric Vehicle Charging Demand Forecasting Techniques," *IEEE Open Journal of Vehicular Technology*, vol. 5, pp. 1348–1373, 2024.

[22] S. Siami-Namini, N. Tavakoli, and A. Siami Namin, "A Comparison of ARIMA and LSTM in Forecasting Time Series," in *2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA)*, Dec. 2018, pp. 1394–1401.

Presenter Biography

Sascha Gohlke is a PhD candidate at the Institute of Data-Centric Software Systems at the Karlsruhe University of Applied Sciences. He holds a Master's degree in Business Information Systems from the Karlsruhe University of Applied Sciences in Germany and a Master's degree in Computer Sciences from the Linnaeus University in Sweden.