Data-Driven Load Profile Forecasting for EV Charging Stations Leveraging Spatial Dependency Modeling

Babak Ravanbach
Energy Systems Technology
DLR Institute of Networked Energy Systems
Oldenburg, Germany
babak.ravanbach@dlr.de

Hrushikesh Mantri
Energy Systems Technology
DLR Institute of Networked Energy Systems
Oldenburg, Germany
hrushikesh.mantri@dlr.de

Chaimaa Essayeh
Department of Engineering
Nottingham Trent University
Nottingham, UK
chaimaa.essayeh@ntu.ac.uk

Samuel Matias
Centre for NEW Energy Technologies
EDP – Energias de Portugal
Porto, Portugal
samuel.matias@edp.pt

***Abstract*—Accurate forecasting of energy demand at electric vehicle charging stations is critical for efficient energy management, reducing grid stress, and supporting the expansion of sustainable transport infrastructure. In this study, we introduce a novel forecasting approach that leverages spatial dependencies across a network of charging stations to improve the predictions of energy demand. Using EMOTION data from the city of Galantina, Italy, our model integrates information from nearby stations, capturing local demand patterns and spatial correlations through the use of two advanced recurrent neural network models, Long Short-Term Memory network (LSTM) and Graph Convolutional Long Short Term Memory (GCLSTM). The model is designed to predict energy demand by learning dynamic interdependencies within the city's charging network, where each station's demand is influenced by spatially adjacent stations and its distance to the city center. The LSTM model, which treats charging stations as independent entities, achieved an overall Mean Absolute Error (MAE) of 0.34, while GCLSTM reached 0.45 MAE. However, GCLSTM outperformed LSTM for stations located near other charging stations, demonstrating its ability to take advantage of spatial dependencies. This suggests that geo-based prediction models are particularly beneficial for dense urban areas with multiple charging stations, such as city centers.**

***Index Terms*—Energy Demand Forecasting, Electric Vehicle Charging Stations, Artificial Neural Networks (ANN), Spatial Dependencies, Urban Mobility.**

I. INTRODUCTION

The global race to achieve net zero emissions has intensified in recent years, with governments and industries implementing ambitious policies to mitigate climate change [1]. A critical component of this transition is the de-carbonization of the transport sector, as it remains one of the main contributors to greenhouse gas emissions [1]. Rapid adoption of electric vehicles presents a promising path toward reducing carbon footprints; however, integrating the demand for charging of electric vehicles into urban energy systems poses significant challenges. Accurate prediction of electric vehicle load profiles is essential to ensure grid stability, optimize charging infrastructure, and enable efficient energy management.

Several studies adopt linear prediction approaches such as ARIMA and SARIMA for predicting electric vehicle (EV) charging demand. In [2], XGboost, LightGBM, SARIMAX, LSTM and Bi-LSTM were applied to EV charging data with hourly resolution to predict 48 hours ahead, all models achieved better results with nearby mean absolute percentage error. Recently, considerable attention has been directed towards neural networks as in [3]. Neural networks have been effective for short term load forecasting as compared to statistical methods, such as regression and ARIMA models. This is due to neural networks ability to handle non-linear data in short term load variations, which are common is load usage data. Unlike traditional methods that rely on long historical data continuity and may not be able to capture volatile data, neural networks learn from the data, improve their forecasts based on real time inputs. Additionally, hybrid models [4] [5] that combine both spatial and time-based data are also being effective to improve forecasting accuracy [6].

Exploring spatial data dependency over the time-based data helps analyzing how it is affecting the prediction results and how important it is to use the spatial data in future aspects. In this work [7] a combined usage of spatial and temporal data using MSTEM (Multiscale Spatio-Temporal Enhanced method) helped in accurately predicting the load profile.

979-8-3315-1278-1/25/$31.00 ©2025 IEEE

Exploring not only the general time series data of EV charging stations, but also their spatial dependency is important in the prediction approach. Therefore, in this work we applied both LSTM and GCLSTM on the EV charging data; where LSTM is using temporal dependency and GCLSTM trying to capture spatial dependency. In addition, the GCLSTM result is compared with LSTM to see how well spatial dependency affects the predictions for the experimented dataset.

In this paper, we present a novel approach for forecasting energy demand at an individual EV charging station by leveraging spatial dependencies across interconnected network of stations. Our contributions are threefold:

- Introducing a data integration technique that enriches EMOTION data with spatial information from charging stations, enabling the model to learn and predict demand with greater contextual awareness.
- Developing deep neural network models specifically designed to capture and utilize spatial correlations in energy demand between charging stations.
- Demonstrating the effectiveness of our approach through comprehensive experiments.

The remainder of the paper is organized as follows. Section II details the data pre-processing steps, including cleaning and feature generation, followed by Section III providing a description of the predictive models and the hyperparameter tuning approach used to optimize performance. Section IV analyzes the forecasting results and provides a discussion on the key findings. Finally, the paper concludes with a summary of insights and potential directions for future research.

II. METHODOLOGY

Predicting load profiles at EV charging stations consists of several integrated steps. Initially, data is collected from the participating charging station manufacturer and operator EMOTION SRL, a pre-processing follows to clean and standardize the data. Feature engineering is then conducted to extract relevant features that significantly impact charging behaviors and load profiles, such as temporal and spatial data.

A. Dataset description

The dataset utilized in the paper is sourced from EMOTION SRL, part of the EU's DriVe2X project, which focuses on technologies for mass EV adoption. The data encompasses records from 133 electric vehicle charging stations across different Italian cities. The City of Galantina (Fig. 1) was selected due to the availability of more charging events compared to other cities.

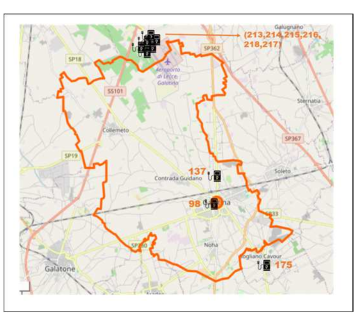

Fig. 1. Locations of charging station, City of Galantina, Italy

Its dataset contains 63,840 slow charging events, recorded from July 26, 2023 to April 17, 2024, measured at 9 charging stations, with each station providing up to 22 kW. These records are not continuous, but event-based, capturing each charging occurrence with precise timestamps and locations defined by latitude and longitude. Every record has a timestamp, longitude and latitude of the charger, a charger identifier and consumed energy ($energy_{cons}$). This spatially spread dataset allows for in-depth spatio-temporal analysis, crucial for leveraging spatial dependencies to improve energy demand predictions.

B. Data cleaning

The cleaning process involves filtering out irrelevant data, focusing only on essential attributes such as timestamps, consumed energy, and identifiers for each charging station. To prepare the dataset for analysis, several cleaning and filtering steps were conducted to ensure reliability and consistency. First, missing values were addressed by removing them. The timestamp column was converted to a proper date-time format, enabling time-based operations such as resampling, aggregation, and filtering. Data consistency was also verified by removing duplicate records to eliminate redundant data points that could distort the analysis. Additionally, geographical coordinates (latitude and longitude) for each station were validated to check for any inaccuracies or mismatched locations.

C. Timestamp resampling

Since the records were event-based, the raw dataset consisted of random time stamps, which vary from 1-minute to 5-minute intervals. A timestamp resampling of the data was performed. The process involves calculating the charging power p_{ch} (kW) of each charging event according to eq. 1 ([8]), and then aggregating charging events into a uniform time interval, typically hourly, which helped smoothing out anomalies and highlighting the general trends in power usage.

$$p_{ch} = \frac{energy_{cons}}{time_{plug-out} - time_{plug-in}} \tag{1}$$

Special attention was given to handling periods of inactivity at charging stations by assigning them a zero power-value, ensuring that the dataset accurately reflects the real-world operation of the stations and maintains continuity in the data [9].

D. Feature generation

The constructed load profiles exhibit sudden peaks and periods of inactivity, making it hard for a prediction model to capture the non-stationary volatile behaviour of the pattern. STL decomposition (Seasonal-Trend Decomposition using Loess) breaks the time-series into three components: trend $T(t)$, seasonality $S(t)$, and residuals $R(t)$ [10], making it effective for handling noisy, non-stationary data. This decomposition is crucial for understanding how different factors such as time of day, day of the week, and seasonal changes impact the load at charging stations, and detect non-stationary patterns of the data. This decomposition reveals long-term trends, periodic behaviors, and random fluctuations in power consumption, providing valuable information on usage patterns.

$$P_{ch}(t) = T(t) + S(t) + R(t) \qquad (2)$$

For this dataset with consideration of multiple seasonality, the seasonal component is divided into daily ($S_{daily}(t)$) and weekly ($S_{weekly}(t)$) components, resulting in:

$$Pch(t) = T(t) + S_{daily}(t) + S_{weekly}(t) + R(t) \qquad (3)$$

Figure 2 highlights the power usage decomposition of *station 98* (for a selected period of two weeks).

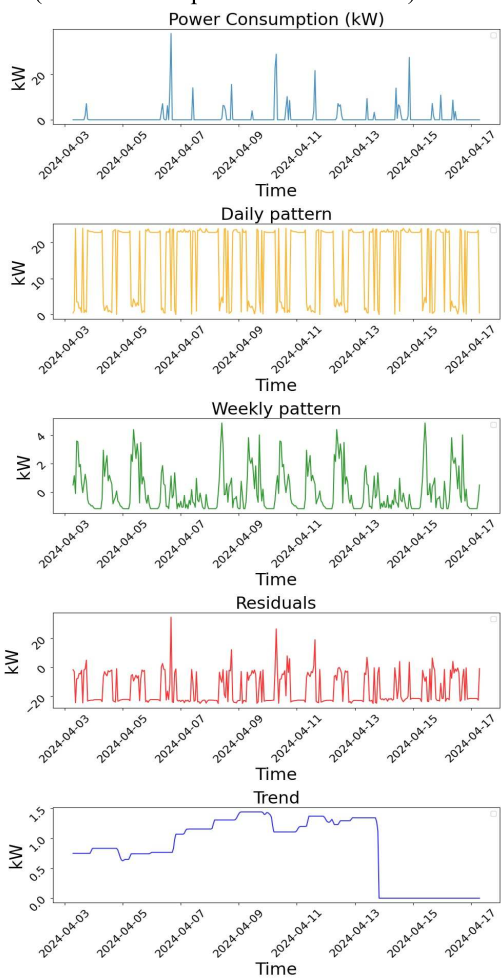

Fig. 2. STL decomposition of power usage of station 98

The increasing trend during certain periods suggests growing station utilization or seasonal influence. Residuals indicate that while seasonality and trend capture much of the variability, occasional unpredictable spikes remain, underscoring the need for robust forecasting methods to address irregular usage patterns. Predictions were made on residuals ($\hat{R}(t)$), later combined with trend and seasonal components to produce final power consumption predictions ($\hat{P}_{ch}(t)$), as explained in [11].

$$\hat{P}ch(t) = T(t) + S_{daily}(t) + S_{weekly}(t) + \hat{R}(t) \qquad (4)$$

Additionally, feature engineering collected spatial input variables to boost model accuracy. This included distance from the station to the city center and inter-station distances.

III. PREDICTION MODELS

Two predictive models, namely, Long-Short Term Memory (LSTM) and Graph Convolutional LSTM (GCLSTM) are implemented and used as a basis for the companions of prediction results for this project.

A. Models

LSTM [12] was used for its well-known effectiveness with sequential data, aimed at capturing longer-term dependencies and variations in EV charging demand. GCLSTM [13] was used for incorporating both temporal and spatial dependencies, which are crucial for this application due to the geographic spread of charging stations.

The input features for the predictive models were carefully structured to capture relevant temporal and spatial dependencies. For LSTM model, the input features included residual power consumption and distance to the city center, represented as: $X_{LSTM}(t) = [R(t), D]$ where $R(t)$ denotes the residual power consumption at time t and D represents the distance to the city center. For the GCLSTM model, additional spatial features, such as geographical coordinates, were incorporated to capture spatial dependencies across charging stations. The input feature structure for the GCLSTM model is given by: $X_{GCLSTM}(t) = [R(t), D, G]$ where G denotes the geographical coordinates of the stations. For the temporal features, each input sequence consisted of a 24-element vector representing one day's worth of data, paired with a 24-hour target representing the next day.

B. Hyperparameter optimization

Hyperparameter optimization is a crucial step in model development because the performance of a model heavily depends on chosen hyperparameters. These parameters, such as learning rates and number of hidden layers, are not learned during training but must be set before training begins. Table I provides a short description of the selected hyperparameters used in the optimization process.

TABLE I. SELECTED HYPERPARAMETERS

Parameters	Description
Hidden units	Number of processors in each layer, higher means more complex information handling
Hidden layers	Stacked processing levels; more layers can capture more complex patterns.
Learning Rate	The parameter that adjusts the step size at each iteration while moving toward a minimum of a loss function
K	Number of neighbouring nodes considered in graph models; larger values increase the area of influence.
Dropout rate	Percentage of processors randomly ignored during training to prevent over-relying on certain paths.
L2 regularization	Technique to keep the model simple and prevent it from being overly precise on training data alone.

To this end, we used and implemented an optimization framework called Optuna [14] for the hyperparameter optimization, which employs a Bayesian optimization technique such as Tree-structured Parzen estimators. It is more efficient than grid or random search, as its dynamic pruning of underperforming trails significantly reduces computational overhead, making it ideal for resource-intensive tasks.

C. Training, Validation, and Testing

The dataset is split into training, validation, and testing sets to optimize the training process and accurately assess the models' performance. The training set spans from July 26, 2023, to January 9, 2024, the validation set from January 9, 2024, to February 27, and the test set from February 27 to April 17, 2024. This split is designed to mirror real-world conditions as closely as possible, allowing for effective training and robust validation of the models. During the training and validation stage, time series cross-validation (TSCV) is applied to ensure robust performance while maintaining temporal integrity. During the training stage, TSCV splits data into sequential training and validation sets through a forward chaining approach. The model is trained on past data and validated on future data on each fold, mimicking real-world scenarios. This method ensures the model learns patterns according to real-world scenarios. During hyperparameter optimization, TSCV evaluates various configurations, such as learning rate.

D. Evaluation Metric

The Mean Absolute Error (MAE) quantifies the average magnitude of errors in predictions, irrespective of their direction. It is defined as:

$$\mathrm{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i|$$

where y_i represents the actual value, $\hat{y}_i$ is the predicted value, and n is the total number of observations. MAE provides an intuitive interpretation of prediction accuracy in the same units as the target variable.

IV. RESULTS

A. Hyperparameter-tuning results

Table II demonstrates the optimal combination of hyperparameter optimization for each model that produced the minimum MAE. The smallest MAE is achieved when LSTM uses the same dropout and L2 regularization but less hidden units and learning rate than GCLSTM. GCLSTM has an additional parameter called K, which represents neighbouring nodes in the GCLSTM architecture. In this optimization, it aggregates 3 neighbouring nodes to achieve the MAE mentioned in the table.

TABLE II. OPTIMAL COMBINATION OF HYPERPARAMETERS

Parameters	**LSTM**	**GCLSTM**
Hidden units	112	128
Hidden layers	2	1
Learning Rate	0.000946	0.0095
Dropout rate	0.30	0.30
L2 regularization	0.20	0.20

B. Forecasting results

The performance of LSTM and GCLSTM models across stations and dates—2024-03-25 (Monday), 2024-03-15 (Friday), and 2024-03-30 (Saturday) are listed in Table III. The MAE values highlight the strengths and limitations of each approach in predicting EV station power consumption. LSTM consistently outperforms the other model, achieving lower errors across most stations and dates by effectively capturing temporal dependencies. GCLSTM performs competitively, excelling in stations with correlated power consumption patterns by leveraging spatial relationships, but it struggles slightly with isolated anomalies, such as sharp peaks in consumption, due to predicting all stations at once. Between the two models, higher MAEs are observed in stations with higher peak usage, reflecting better in capturing larger trends. Temporal variability further reveals that LSTM maintains stable performance across weekdays and weekends, whereas GCLSTM is sensitive to noise in spatial dependencies during busy periods. These results emphasize the importance of tailoring models to the temporal and spatial characteristics of EV station data, with LSTM providing robust generalizability and GCLSTM might help in scenarios requiring spatially-aware predictions. But with an extra dataset and with huge co-relation with adjacent stations, GCLSTM might outperform LSTM.

TABLE III. COMPARISON OF THE PERFORMANCE OF EACH STATION BETWEEN LSTM AND GCLSTM

Stations	**2024-03-25 (Mon)**		**2024-03-15 (Fri)**		**2024-03-30 (Sat)**	
	LSTM	***GCLSTM***	***LSTM***	***GCLSTM***	***LSTM***	***GCLSTM***
98	0.297	0.441	0.389	0.563	0.306	0.378
137	0.263	0.247	0.652	0.935	0.312	0.303
175	0.282	0.277	0.216	0.331	0.228	0.251
213	0.193	0.256	0.219	0.230	0.254	0.258
214	0.274	0.256	0.265	0.260	0.348	0.237
215	0.169	0.259	0.249	0.260	0.419	0.601
216	0.250	0.208	0.274	0.247	0.222	0.340
217	0.892	1.878	0.332	0.489	0.264	0.349
218	0.407	0.829	0.688	0.721	0.630	0.681

Low error in model performance, **High error** in model performance

Although the exhibited samples show that the models can track power consumption, they also display the weakness of the models, which is the inability to accurately predict zero values. During the weekend (Saturday), station 214 and 215, which are near the airport, have parallel activity across stations. Still, GCLSTM predicted power consumption peaks without any large errors, suggesting that GCLSTM helps in predicting a network of stations using their inter-station relations. For all the three days we can notice the repeated power consumption at station 98 during midday which is the reason LSTM is able to predict well because of the repeating patterns.

Figure 3 illustrates the performance of both models, LSTM and GCLSTM, on different days (weekends, weekdays, and random high variability days, Monday in this case) for selected charging stations.

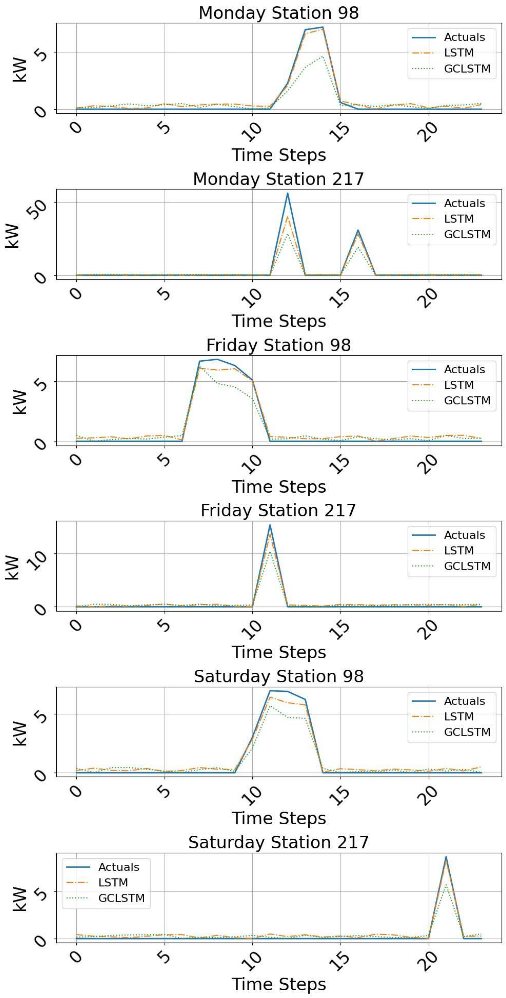

Fig. 3. Forecasting results for the 24 hours for selected charging sites for 2024-03-25 (Monday), 2024-03-15 (Friday), and 2024-03-30 (Saturday)

V. CONCLUSION AND FUTURE WORK

This paper studied the use of advanced models to predict power use at EV charging stations by testing models like LSTM and GCLSTM. The results show that the LSTM model is more accurate and reliable, as it effectively captures long-term patterns in power use. The GCLSTM model also performed well, as it can account for both location and time-based patterns, which is especially useful for station-dense networks, hence make it ideal for interconnected networks of stations where correlation between stations is crucial. The spatial dependency was considered for this project in order to understand the influence of such type of data in the prediction process as an additional input, compared to conventional state-of-the-art models that take only time-dependent data as input. Our modeling methodology and architecture allows for two sets of inputs, and considering their interrelation in our prediction approach, it sets the ground for inclusion of more spatial data with potentially higher correlation between stations and point of interest on the map in future work coming from real-world environments. For example, implementing the model in smart city infrastructure pilots that can provide the prediction tool with more dynamic real-time data features such as traffic, building occupancy, weather, EV location, etc. Future research should also consider approaches to reduce computational cost and increasing speed through methods like model pruning, quantization, and edge computing. Developing hybrid and ensemble models, which combine the strengths of different approaches, could lead to improved reliability and performance. By pursuing these research directions, the field can move closer to creating effective and intelligent energy management solutions that support the rising demand for EV infrastructure and contribute to sustainable energy systems.

ACKNOWLEDGMENT

The authors would like to thank the funding for DriVe2X research and innovation project from the European Commission and the UKRI, with grant numbers 101056934 and 10055673, respectively.

REFERENCES

[1] I. E. A. (IEA), “Global ev outlook 2021: Accelerating ambitions despite the pandemic,” 2021, available at: https://www.iea.org/reports/global-evoutlook-2021.

[2] M. Clerc, “Short-term forecasting of ev charging stations power consumption at distribution scale,” Master’s thesis, KTH Royal Institute of Technology, Stockholm, Sweden, 6 2021, master’s Programme in Environmental Engineering and Sustainable Infrastructure.

[3] D. Panahi, S. Deilami, M. A. Masoum, and S. M. Islam, “Forecasting plug-in electric vehicles load profile using artificial neural networks,” *IEEE Xplore*, 1 2015, dOI: 10.1109/ACCESS.2015.2438159. [Online]. Available: https://ieeexplore.ieee.org/document/2438159

[4] Y. Wang, K. Qian, M. Han, and X. Peng, “Short-term load forecasting of electric vehicle charging stations based on hybrid machine learning model,” *Energy*, vol. 191, p. 116548, 2020.

[5] Z. Wang, C. Liu, and Y. Zha, “Load forecasting for electric vehicle charging stations using a data-driven approach,” *IEEE Transactions on Transportation Electrification*, vol. 6, no. 3, pp. 123–133, 2020.

[6] M. A. Hammad, B. Jereb, B. Rosi, and D. Dragan, “Methods and models for electric load forecasting: A comprehensive review,” *Logistics & Sustainable Transport*, vol. 11, no. 1, pp. 51–76, 2020.

[7] Z. Zhang, J. Hao, W. Zhao, Y. Liu, Y. Huang, and X. Luo, “Multiscale spatio-temporal enhanced short-term load forecasting of electric vehicle charging stations,” *arXiv preprint arXiv:2405.19053*, 2023.

[8] Y. Ali, F. Monsuur, C. Morton, and C. K. Man, “Modelling electric vehicle charging events using a random parameter duration approach,” *Travel Behaviour and Society*, vol. 39, p. 100989, January 2025, available online 20 January 2025. [Online]. Available: https://doi.org/10.1016/j.tbs.2025.100989.

[9] M. Alizadeh, A. Scaglione, and A. Davoudi, “Optimal real-time electric vehicle charging for renewable energy integration,” *IEEE Transactions on Smart Grid*, vol. 9, no. 2, pp. 1518–1530, 2018.

[10] S. Seabold and J. Perktold, “statsmodels: Econometric and statistical modeling with python,” 2010. [Online]. Available: https://www.statsmodels.org/stable/index.html

[11] R. Fonteijn, T. Castelijns, M. Grond, P. Nguyen, J. Morren, and H. Slootweg, “Short-term load forecasting on mv/lv transformer level,” in *25th International Conference on Electricity Distribution*. Madrid, Spain: CIRED, Jun. 2019, pp. 1–5, paper No. 559.

[12] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[13] J. Chen, X. Wang, and X. Xu, “Gc-lstm: Graph convolution embedded lstm for dynamic network link prediction,” *Preprint submitted to an undisclosed journal*, 2021, arXiv:1812.04206v2.

[14] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, “Optuna: A nextgeneration hyperparameter optimization framework,” in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019.