Journal of Applied Data Sciences
Vol. 7, No. 1, January 2026, pp. 409-424

TF-EffBiGRU-AttNet: A Novel Deep Learning Framework for Spatio-Temporal Energy Demand Forecasting in Electric Vehicle Charging Networks

S. Prakash[1], S. Aruna Mary[2,*], G. Sudhagar[3], M. Batumalay[4,5]

[1]*Department of EEE, Bharath Institute of Higher Education and Research, Chennai, India*

[2,3]*Department of ECE, Bharath Institute of Higher Education and Research, Chennai, India*

[4]*Faculty of Data Science and IT, INTI International University, 71800 Nilai, N. Sembilan, Malaysia*

[5]*Centre for Data Science and Sustainable Technologies, INTI International University, 71800Nilai, N. Sembilan, Malaysia*

(Received: May 19, 2025; Revised: July 1, 2025; Accepted: October 20, 2025; Available online: December 19, 2025)

Abstract

Electric Vehicle Charging Stations (EVCS) are key enablers of sustainable transportation, yet accurate forecasting of their energy demand remains challenging due to complex spatial-temporal variability. This study introduces a novel hybrid deep learning framework, Two-Fold EfficientNetV2 BiGRU with Attention (TF-EffBiGRU-AttNet), optimized using the Self-Adaptive Hippopotamus Optimization Algorithm (SA-HOA), to enhance prediction accuracy and computational efficiency in EVCS energy demand forecasting. The main objective is to integrate multi-scale spatial learning, bidirectional temporal modeling, and adaptive feature prioritization within a single architecture capable of robust and interpretable forecasting. The model's novelty lies in its dual-fold spatial feature extraction using EfficientNetV2 and dynamic optimization through SA-HOA, which adaptively balances exploration and exploitation during training. Experimental validation on two real-world datasets from Palo Alto and Perth demonstrates that the proposed model consistently outperforms state-of-the-art baselines. For the 7-1 forecasting task, TF-EffBiGRU-AttNet achieved the lowest MAE of 0.012 and RMSE of 0.051 for Palo Alto, and MAE of 0.029 with RMSE of 0.12 for Perth. For the 30-7 task, it achieved MAE of 0.0332, RMSE of 0.1654, and MAPE of 0.20% on Palo Alto, and MAE of 0.0235, RMSE of 0.0824, and MAPE of 0.37% on Perth, outperforming Bi-LSTM and EfficientNet by over 60% in RMSE reduction. Moreover, SA-HOA improved optimization efficiency with a best fitness value of 0.0003 and reduced convergence time to 1.2 seconds, surpassing PSO, GWO, and HOA. These results highlight the framework's ability to capture spatial-seasonal and nonlinear dependencies while maintaining low computational overhead. The findings confirm the model's potential as a robust, adaptive, and scalable solution for intelligent EV energy demand forecasting, supporting smart grid planning and sustainable energy management.

Keywords: Spatio-Temporal Forecasting, Electric Vehicle (EV) Charging Networks, Deep Learning Framework, Efficient BiGRU (EffBiGRU), Attention Mechanism, Process Innovation.

1. Introduction

The rapid adoption of Electric Vehicles (EVs) has significantly increased the demand for intelligent energy management systems in charging infrastructures. The energy demand at EV charging stations exhibits strong spatio-temporal characteristics influenced by factors such as user behavior, traffic conditions, and time-of-day consumption patterns [1]. Accurate spatio-temporal energy demand forecasting is therefore essential for anticipating these variations and optimizing power resource utilization. However, conventional forecasting approaches such as statistical and traditional time-series models often fail to capture the complex interdependencies between spatial and temporal factors [2].

Spatio-temporal forecasting offers several advantages. It enhances grid stability, prevents both overloading and underutilization of charging stations, and enables proactive energy distribution, thereby reducing operational costs and grid stress [3]. Moreover, accurate forecasting facilitates the integration of renewable energy sources, contributing to sustainability and energy efficiency within EV charging networks. As EV adoption continues to grow, advanced

***Corresponding author: S. Aruna Mary (aruna.ece@bharathuniv.ac.in)**
DOI: https://doi.org/10.47738/jads.v7i1.805

forecasting techniques are becoming indispensable for optimizing charging station placement, balancing power loads, and supporting vehicle-to-grid (V2G) interactions [4]. Despite its importance, forecasting energy demand in EV charging networks presents multiple challenges due to the dynamic and nonlinear nature of charging behaviors [5]. Demand variability driven by fluctuations in user activity, weather, and traffic complicates predictive modeling. Spatial correlations among charging stations depend not only on geographical proximity but also on real-time energy demand shifts, necessitating adaptive and data-driven modeling approaches [6]. Traditional models often struggle to address the multiscale and nonlinear dependencies inherent in spatio-temporal data, resulting in limited forecasting precision. In addition, data sparsity, inconsistency, and the computational burden of handling real-time data streams further constrain prediction accuracy and scalability, particularly in large urban networks [7]. Consequently, there is a growing need for robust, scalable, and interpretable learning models capable of efficiently modeling spatio-temporal dependencies while maintaining computational efficiency and data security [8].

Several recent studies have introduced advanced deep learning frameworks to address these challenges. For instance, the Attribute-Augmented Spatio-Temporal Graph Informer (AST-GIN) integrates Graph Convolutional Networks (GCN) and Informer architectures to enhance prediction accuracy and interpretability in transportation systems [9]. Similarly, the Simultaneous Demand Prediction and Planning (SPAP) framework employs multi-source feature learning and iterative optimization to improve charger deployment strategies across cities [10]. Other models incorporate spatial enhancement modules or hierarchical frameworks to capture both short- and long-range dependencies and to enable large-scale forecasting across diverse contexts [11], [12], [13].

Building upon these advancements, this study proposes TF-EffBiGRU-AttNet, a novel deep learning framework for spatio-temporal energy demand forecasting in EV charging networks. The proposed model integrates Transformer-based temporal learning with a Two-Fold EfficientNetV2 backbone for spatial feature extraction, a Bidirectional Gated Recurrent Unit (BiGRU) for sequential dependency modeling, and an Attention Mechanism for adaptive feature prioritization. By combining these modules with a Self-Adaptive Hippopotamus Optimization Algorithm (SA-HOA), the proposed approach aims to improve forecasting accuracy, computational efficiency, and adaptability to dynamic charging environments [14], [15].

2. Literature Review

Wang et al. [16] proposed an Adaptive Spatio-Temporal Graph Recurrent Network (ASTGRN) for short-term electric vehicle (EV) charging demand prediction in smart grids and transportation systems. The model integrates adaptive graph learning with an embedding projection layer to capture complex spatial and temporal dependencies. Validated using GPS trajectory data of EVs, the framework demonstrated superior forecasting accuracy compared to existing models.

Kuang et al. [17] developed CityEVCP, a hypergraph network model that utilizes a Transformer encoder to predict city-scale EV charging demand. By applying graph attention to learn non-pairwise relationships and clustering service areas based on Points of Interest (POIs), the model significantly improved forecasting accuracy and enhanced infrastructure planning for sustainable urban mobility.

Chen et al. [18] introduced a Dynamic Time Warping-based Adaptive Spatio-Temporal Graph Convolutional Network (ASTGCN) to enhance EV charging demand prediction. The model employs a FastDTW-based adjacency matrix to capture spatial correlations and integrates ASTGCN to represent spatio-temporal dependencies. Results show improved precision and stability, enabling grid operators to manage fluctuating EV charging demands more effectively.

Yang et al. [19] designed a model for forecasting EV charging load distribution in a transportation–power coupled network (TPCN). The framework combines Dijkstra's algorithm for optimal route planning with an origin–destination (OD) matrix for trip distribution. Monte Carlo simulations were used to model temporal and spatial variations in charging loads, improving the accuracy of load dynamics simulation and station placement decisions.

Marlin et al. [20] developed a Hierarchical Federated Learning Transformer Network (H-FLTN) that focuses on privacy-preserving EV charging demand forecasting. The model integrates secure aggregation and peer-to-peer sharing mechanisms to enhance privacy and efficiency. Tested on large-scale EV mobility datasets, it demonstrated improved resource allocation, grid stability, and predictive accuracy while maintaining user data confidentiality.

Han and Li [21] proposed the V2AFedEGAT model for EV charging station load prediction, combining Long Short-Term Memory (LSTM) networks with hybrid spatio-temporal attention mechanisms. The framework uses client selection strategies to optimize model updates and ensure data privacy. When evaluated on combined transportation and power distribution datasets, the model achieved high accuracy and sub-second response times, indicating its suitability for real-time applications.

Zhuang et al. [22] established an Adaptive Spatio-Temporal Graph Convolutional Network (ASTGCN) for real-time estimation of EV hosting capacity. The approach integrates probabilistic forecasting and risk analysis to extract spatio-temporal features from real-world EV charging data, improving grid capacity management and reducing demand volatility risks.

Hou et al. [23] proposed a spatio-temporal forecasting model that incorporates an enhanced PageRank algorithm and an attention mechanism to improve EV charging demand prediction. Using data from the Hebei South Network, the model achieved higher prediction accuracy and helped grid operators manage fluctuating charging loads more effectively.

Hu et al. [24] developed a GCN–Transformer hybrid model to improve load forecasting for EV battery swapping stations. This model leverages Graph Convolutional Networks to aggregate spatial information and Transformer networks to capture sequence dynamics. It effectively learned station-specific charging patterns and provided robust forecasting results for complex networked environments.

Chen et al. [25] designed SEDformer, a spatio-temporal forecasting model that combines temporal and spatial encoder blocks, sequence decomposition, self-attention, and channel attention mechanisms to efficiently capture intricate spatio-temporal dependencies. Validated on real-world Palo Alto data, the model demonstrated strong performance for long-term EV charging load prediction and contributed to improved energy management and infrastructure planning. A summary of these existing studies, including their methodological approaches, advantages, and limitations, is presented in table 1.

Table 1. Comparison of Existing Studies

Author(s)	Method	Advantages	Disadvantages
Wang et al. [16]	ASTGRN	Data-driven spatial learning, high accuracy	Sensitive to geographical data variations
Kuang et al. [17]	CityEVCP	Captures area relationships, improves forecasting	Complex hypergraph clustering
Chen et al. [18]	ASTGCN + FastDTW	High spatial-temporal accuracy	High computational cost
Yang et al. [19]	TPCN + Monte Carlo	Considers external factors	Requires refined input data
Marlin et al. [20]	H-FLTN	Privacy-preserving, efficient forecasting	High computational demand
Han & Li [21]	V2AFedEGAT + LSTM	Integrates transport and grid data	Complex model training
Zhuang et al. [22]	ASTGCN	High accuracy, effective real-time estimation	Data-intensive
Hou et al. [23]	PageRank-based	Improves node importance estimation	Sensitive to load fluctuations
Hu et al. [24]	GCN–Transformer	Handles complex spatial-temporal dependencies	Limited scalability
Chen et al. [25]	SEDformer	Strong long-term forecasting performance	Requires high-quality input data

 ISSN 2723-6471

Existing studies on EV charging demand forecasting, including ASTGRN [16], CityEVCP [17], ASTGCN [18], and TPCN-based models [19], primarily rely on deep learning and graph-based architectures to capture spatio-temporal dependencies. While these methods have achieved considerable progress, several key challenges remain unresolved. Most models suffer from high computational complexity, limited scalability for dynamic large-scale datasets, and dependency on pre-defined spatial relationships.

Furthermore, frameworks that integrate transportation and power grid data often encounter privacy concerns and rely heavily on large labeled datasets, which are difficult to obtain in real-world scenarios. To overcome these limitations, there is a need for lightweight and adaptive forecasting models that can balance computational efficiency and predictive accuracy. Future research should emphasize the integration of deep learning with optimization algorithms, incorporate real-time data assimilation and self-learning capabilities, and explore decentralized privacy-preserving methods to ensure secure, scalable, and accurate EV charging demand forecasting.

3. Methodology

The proposed TF-EffBiGRU-AttNet framework is designed to accurately forecast energy demand in electric vehicle (EV) charging networks by integrating spatial, temporal, and attention-based learning with adaptive optimization. The model captures complex spatio-temporal dependencies within high-dimensional datasets and maintains computational efficiency through deep learning and metaheuristic optimization. The framework employs the Two-Fold EfficientNetV2 network for spatial feature extraction, the Bidirectional Gated Recurrent Unit (BiGRU) for temporal pattern learning, and an Attention Mechanism to emphasize important features. Model parameters are refined using the Self-Adaptive Hippopotamus Optimization Algorithm (SA-HOA), ensuring rapid convergence and robust performance. The overall workflow of the proposed framework is presented in figure 1, which outlines the major stages from data preprocessing to the final prediction output.

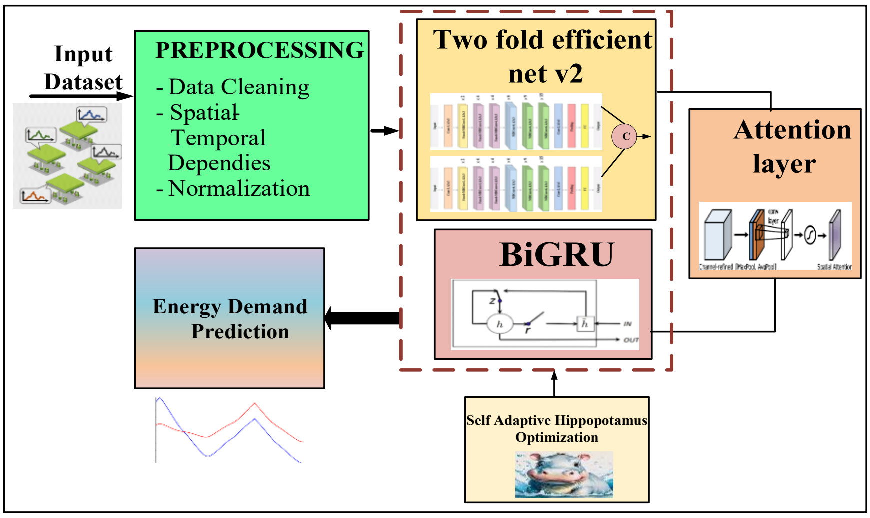

Figure 1. Overall workflow of the proposed TF-EffBiGRU-AttNet model.

3.1. Framework Overview

The methodology begins with data preprocessing, including cleaning, handling missing values, and extracting spatial and temporal dependencies. The processed data are then passed into the Two-Fold EfficientNetV2 module, which captures complex spatial patterns across EV charging stations. The output of this module is refined using an attention layer that highlights critical spatial features. These enhanced spatial representations are subsequently fed into the BiGRU module, which learns temporal correlations and dynamic variations in energy demand over time.

After temporal learning, the model parameters are optimized through SA-HOA to improve training stability and prediction accuracy. The optimized framework produces the final energy demand forecast for the EV charging network, capable of adapting to changes in user behavior, traffic density, and environmental factors.

3.2. Dataset Description

This study uses two benchmark datasets, PERTH and PALO ALTO, which record electric vehicle charging activity between July 2011 and December 2020. The datasets contain spatial and temporal information from stations across Perth, Kinross, and Palo Alto. They include patterns influenced by seasonal variations, time-of-day charging behavior, and traffic conditions. This variety of data enables the proposed model to learn diverse energy consumption behaviors and improves its generalization to different urban and geographical settings.

3.3. Data Preprocessing

Preprocessing is a crucial step to ensure data consistency and quality before model training. Missing data, often caused by sensor failures or transmission interruptions, are reconstructed using a weighted average of previous values as expressed in Equation (1). Outliers that deviate more than three standard deviations from the mean are identified through Z-score analysis and replaced using interpolation methods. This process reduces noise and prevents anomalies from affecting model learning.

$$x_t = \frac{1}{k}\sum_{d=1}^{k} x_{d,t} \tag{1}$$

Spatial relationships among charging stations are represented as an undirected graph $G = (V, E)$, with nodes representing the charging stations and edges representing weighted spatial dependencies. The adjacency matrix $A$is dynamically updated according to variations in traffic flow and station proximity, allowing the model to capture real-time spatial dependencies across the network.

$$G = (V, E) \tag{2}$$

Temporal dependencies are modeled using features such as time of day, day of the week, and seasonal indicators. Short-term dependencies are represented through autoregressive time steps, while long-term dependencies are captured through periodic temporal sequences. All input features are normalized using Min–Max normalization, ensuring stable convergence during training and balanced contribution from each feature, as shown in Equation (3).

$$x' = \frac{x - x_{min}}{x_{max} - x_{min}} \tag{3}$$

The datasets are divided into training, validation, and testing subsets with proportions of 70%, 15%, and 15%. Temporal order is preserved to maintain the sequential integrity of the data and prevent information leakage across the subsets.

3.4. Spatial Feature Extraction using Two-Fold EfficientNetV2

The Two-Fold EfficientNetV2 network is employed to extract spatial dependencies between EV charging stations. This enhanced version of EfficientNetV2 uses compound scaling that adjusts depth, width, and resolution simultaneously, allowing efficient learning with minimal computational cost. The two-fold configuration improves spatial representation by capturing both localized and global spatial dependencies.

The first fold processes fine-grained features, focusing on short-range spatial dependencies among nearby stations. The second fold focuses on broader regional correlations, allowing the model to learn large-scale energy consumption trends. The extracted spatial features are then passed through fully connected layers, producing an integrated representation that combines both local and regional spatial information. This spatial foundation provides critical input for the subsequent temporal modeling process.

3.5. Temporal Modeling with BiGRU and Attention Mechanism

Temporal dependencies are captured through a Bidirectional Gated Recurrent Unit (BiGRU), which processes information in both forward and backward directions. This bidirectional design allows the model to learn how previous and future charging behaviors influence current energy demand. The BiGRU captures both short-term fluctuations and long-term patterns, offering improved learning efficiency compared to traditional RNNs or LSTMs (see figure 2).

An Attention Mechanism is incorporated after the BiGRU layer to improve model interpretability and focus. The attention component assigns adaptive weights to time steps and features that have the most influence on demand variation. This allows the model to prioritize crucial temporal segments and ignore less relevant ones. The combination of BiGRU and attention strengthens the model's ability to detect essential spatio-temporal patterns while maintaining high prediction accuracy.

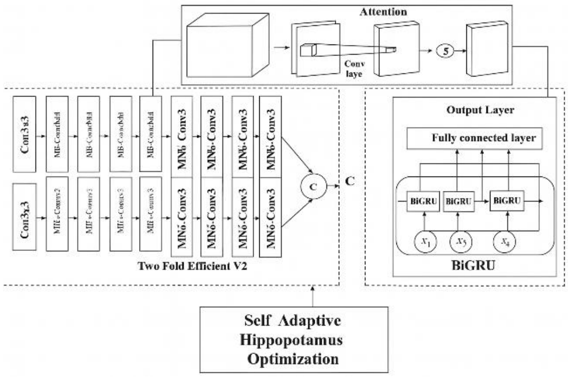

Figure 2. Architecture of the Two-Fold EfficientNetV2 BiGRU with integrated Attention Mechanism.

3.6. Optimization using Self-Adaptive Hippopotamus Optimization Algorithm (SA-HOA)

The Self-Adaptive Hippopotamus Optimization Algorithm (SA-HOA) is applied to optimize model parameters and improve convergence. This metaheuristic algorithm is inspired by the natural behaviors of hippopotamuses, including exploration, cooperation, and defense mechanisms. It balances global exploration and local exploitation through an adaptive strategy, ensuring faster convergence and improved solution quality. In the initialization stage, a population matrix $Y_i$is generated to represent candidate solutions within the search space. Each position $y_{ij}$is initialized using the equation

$$y_{ij} = lb_j + r(ub_j - lb_j) \tag{4}$$

with $lb_j$and ub_jrepresenting the lower and upper limits of the search boundaries, and $r$being a random number between 0 and 1. During the exploration phase, male hippopotamuses compete to identify dominant solutions, and their positions are updated according to the fitness of these dominant individuals. The equation for position updating is expressed as

$$y_{ij}^{mhippo} = y_{ij} + x_{ij}(D_{hippo} - I_{1x_{ij}}) \tag{5}$$

The exploitation phase focuses on refining the best-found regions using the cooperative behavior of females and juveniles. The exploration factor $T = e^{-l/t}$decreases gradually, allowing the algorithm to shift from exploration to exploitation. This adaptive adjustment enhances the balance between global search and local refinement.

By dynamically modifying its internal parameters, SA-HOA effectively accelerates convergence, reduces computational time, and improves forecasting accuracy. The integration of SA-HOA into the training process of TF-EffBiGRU-AttNet ensures optimal parameter selection, making the framework robust and efficient across diverse datasets.

The proposed TF-EffBiGRU-AttNet model combines multi-scale spatial learning through Two-Fold EfficientNetV2, temporal sequence modeling via BiGRU, and adaptive attention-based feature prioritization. The incorporation of SA-HOA further enhances optimization by achieving stable and rapid convergence. Together, these components create a unified deep learning framework capable of delivering accurate, scalable, and computationally efficient spatio-temporal energy demand forecasts for electric vehicle charging networks.

4. Result and Discussion

The proposed TF-EffBiGRU-AttNet model demonstrates superior performance compared to conventional and state-of-the-art baseline models in terms of both accuracy and generalization capability. Extensive experiments were conducted on multiple datasets to evaluate the robustness, scalability, and adaptability of the model under different forecasting tasks. Two forecasting settings were tested: the 7-1 task, which predicts one future step based on seven historical time steps, and the 30-7 task, which predicts seven future steps using 30 historical observations. In addition to performance comparisons, an ablation study was performed to analyze the contribution of individual components of the model, including the attention mechanism, spatial encoder, and temporal sequence learner.

The model evaluation used three standard statistical metrics: Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Mean Absolute Percentage Error (MAPE). These metrics measure prediction deviation, overall error dispersion, and percentage-based forecasting accuracy, respectively. Lower values of MAE, RMSE, and MAPE indicate better forecasting performance and higher model precision.

4.1. Performance Comparison on Task 7-1

The results of the 7-1 forecasting task are summarized in table 2 and visualized in figure 3, 4, and 5. This task involves predicting one future energy demand step based on the previous seven time steps. The evaluation was carried out on two real-world datasets, Palo Alto and Perth, to assess the proposed model's ability to generalize across different geographic and behavioral contexts.

Table 2 provides a comparative analysis of five models—TF-EffBiGRU-AttNet (Proposed), EfficientNet, Bi-GRU, GRU, and Bi-LSTM—using three key performance metrics: Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Mean Absolute Percentage Error (MAPE). These metrics collectively measure prediction accuracy, error magnitude, and percentage-based deviation.

Table 2. Comparison of predictive performance of different models for Task 7-1.

Task 7-1	**Palo Alto**			**Perth**		
	MAE	**RMSE**	**MAPE**	**MAE**	**RMSE**	**MAPE**
TF-EffBiGRU-AttNet (Proposed)	0.012	0.051	0.07%	0.029	0.12	0.45%
EfficientNet	0.140	0.70	0.90%	0.070	0.36	1.20%
Bi-GRU	0.190	1.00	1.20%	0.100	0.50	1.60%
GRU	0.180	0.95	1.13%	0.095	0.48	1.50%
Bi-LSTM	0.072	0.51	0.43%	0.045	0.31	0.89%

As shown in table 2, the proposed TF-EffBiGRU-AttNet consistently achieves the lowest values for all error metrics across both datasets, indicating superior performance in forecasting accuracy and stability. On the Palo Alto dataset, the model achieves a remarkably low MAE of 0.012 and RMSE of 0.051, outperforming all baseline models by a large margin. This suggests that the proposed model can effectively capture both short-term and long-term temporal dependencies. The Perth dataset exhibits slightly higher errors due to more irregular spatial-temporal dynamics, reflecting the city's variable charging behaviors influenced by urban layout and demand density. Nevertheless, the model still demonstrates clear superiority, achieving MAE and RMSE values of 0.029 and 0.12, respectively—significantly lower than competing methods.

Figure 3 illustrates the MAE comparison between the five models for both datasets. The MAE metric quantifies the average magnitude of forecasting errors, providing a direct measure of prediction accuracy. In Palo Alto, TF-EffBiGRU-AttNet yields the smallest MAE, close to zero, which reflects its ability to minimize prediction deviation. In contrast, the Bi-GRU and GRU models exhibit much higher MAE values, indicating larger deviations and instability in capturing demand patterns. The Bi-LSTM performs moderately well but still lags behind the proposed model, suggesting that unidirectional recurrence and the absence of efficient feature selection reduce its predictive strength.

In Perth, a similar pattern is observed: the proposed model maintains the lowest MAE, followed by Bi-LSTM and EfficientNet, while Bi-GRU and GRU show the highest errors. This consistent performance across two geographically distinct datasets validates the robustness and adaptability of TF-EffBiGRU-AttNet in handling different user and network behaviors.

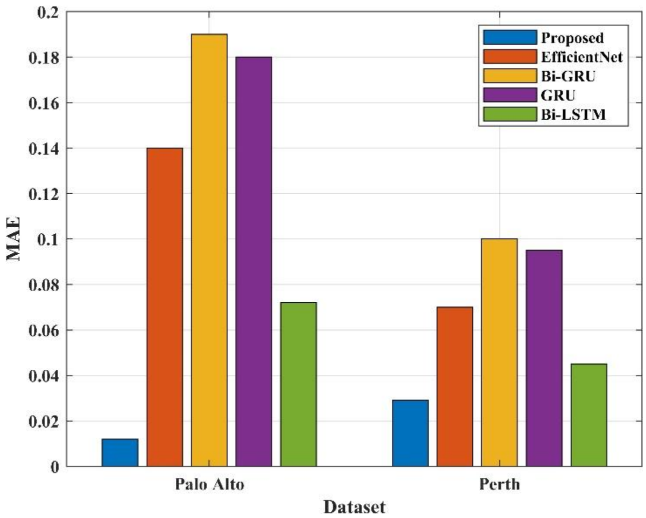

Figure 3. MAE comparison of different models on Palo Alto and Perth datasets.

Figure 4 compares the RMSE values of all models for both datasets. RMSE provides insight into the model's ability to handle large prediction deviations and is particularly sensitive to significant forecasting errors.

The proposed TF-EffBiGRU-AttNet achieves the lowest RMSE in both datasets, confirming its capability to suppress large deviations in forecasting outputs. On Palo Alto, the proposed model's RMSE is 0.051, whereas the second-best model, Bi-LSTM, records 0.51, which is approximately ten times higher. This drastic improvement emphasizes the model's strong learning efficiency in complex temporal spaces.

In the Perth dataset, TF-EffBiGRU-AttNet again delivers the lowest RMSE (0.12), followed by Bi-LSTM (0.31) and EfficientNet (0.36). The results demonstrate that incorporating both bidirectional temporal learning and adaptive attention allows the model to generalize effectively across datasets with differing spatial variability and noise characteristics.

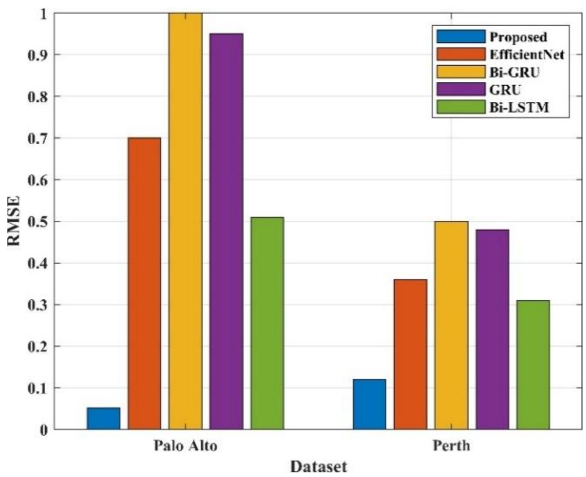

Figure 4. RMSE comparison of different models on Palo Alto and Perth datasets.

Figure 5 shows the MAPE comparison for all models, offering a percentage-based evaluation of forecasting accuracy relative to actual demand. Lower MAPE values indicate better model performance and higher reliability.

Across both datasets, TF-EffBiGRU-AttNet achieves the smallest MAPE—0.07% for Palo Alto and 0.45% for Perth—outperforming all competing models. EfficientNet, Bi-GRU, and GRU produce noticeably higher MAPE values, indicating less stability in predicting the fluctuating nature of EV charging demand. Bi-LSTM performs better than these traditional models but remains inferior to the proposed approach.

The slightly higher MAPE for Perth again reflects the increased variability and complexity in this dataset, which includes more dispersed charging stations and less consistent user demand patterns. Despite this, the proposed model maintains significantly lower percentage errors, confirming its superior adaptability and precision.

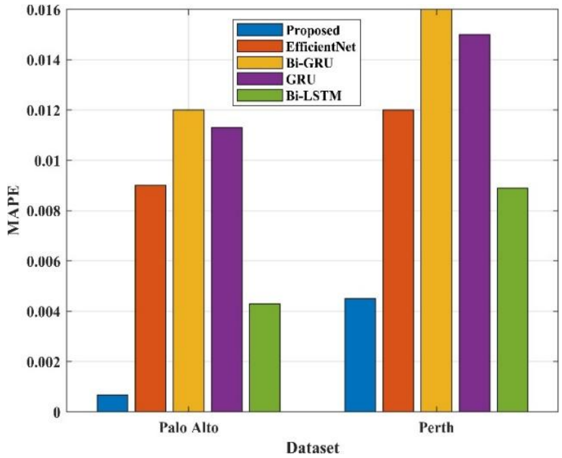

Figure 5. MAPE performance comparison of different models on Palo Alto and Perth datasets.

The experimental findings from both datasets clearly demonstrate that TF-EffBiGRU-AttNet consistently surpasses the performance of EfficientNet, Bi-GRU, GRU, and Bi-LSTM across all three metrics—MAE, RMSE, and MAPE. The model's superior results stem from its hybrid architecture, which combines the strengths of convolutional spatial feature extraction, bidirectional temporal learning, and an adaptive attention mechanism that dynamically focuses on the most relevant temporal and spatial features.

The lower error rates in Palo Alto suggest that the model performs best in regions with stable and structured energy demand patterns, while the slightly higher errors in Perth reflect challenges posed by heterogeneous data distributions and higher variability in charging behavior. Nonetheless, the proposed model's consistent superiority across both contexts demonstrates its capability to generalize effectively to complex, real-world EV charging environments.

Overall, the results from table 2 and figures 3–5 validate that TF-EffBiGRU-AttNet not only minimizes prediction errors but also maintains stability across varying data conditions, making it a robust and reliable framework for intelligent energy demand forecasting in electric vehicle charging networks.

4.2. Performance Comparison on Task 30-7

The 30-7 forecasting task evaluates the model's ability to perform multi-step energy demand prediction, in which seven future time steps are forecasted using the preceding thirty historical time steps. This task is more challenging than the 7-1 setting, as errors can accumulate across multiple forecasting horizons. Therefore, it serves as a strong indicator of a model's long-term generalization and stability.

Table 3 summarizes the performance comparison among the five models—TF-EffBiGRU-AttNet (Proposed), EfficientNet, Bi-GRU, GRU, and Bi-LSTM—based on three key evaluation metrics: MAE, RMSE, and MAPE. The results are reported for both Palo Alto and Perth datasets, providing a comprehensive assessment across diverse urban contexts.

Table 3. Comparison of predictive performance of different models for Task 30-7.

Task 30-7	**Palo Alto**			**Perth**		
	MAE	**RMSE**	**MAPE**	**MAE**	**RMSE**	**MAPE**
TF-EffBiGRU-AttNet (Proposed)	0.0332	0.1654	0.20%	0.0235	0.0824	0.37%
EfficientNet	0.14	0.72	0.84%	0.0968	0.4675	1.64%
Bi-GRU	0.2054	1.2377	1.45%	0.1115	0.5539	2.32%

GRU	0.18	0.96	1.13%	0.095	0.49	1.50%
Bi-LSTM	0.0937	0.4861	0.44%	0.04	0.2697	0.71%

The results show that the TF-EffBiGRU-AttNet model achieves the best overall performance across all evaluation metrics. On the Palo Alto dataset, it reaches a MAE of 0.0332 and RMSE of 0.1654, which are significantly lower than all other baseline models. Even the well-performing Bi-LSTM model shows higher error values, indicating the superiority of the proposed architecture in capturing long-range temporal dependencies. The MAPE value of 0.20% further confirms that the model produces highly accurate percentage-based forecasts with minimal deviation from actual energy demand.

In the Perth dataset, the proposed model maintains its advantage, achieving the lowest MAE (0.0235), RMSE (0.0824), and MAPE (0.37%). Despite the higher temporal variability and less structured demand patterns in this dataset, the TF-EffBiGRU-AttNet continues to outperform all other models. The results suggest that the proposed hybrid structure, combining EfficientNet-based spatial encoding, bidirectional recurrent learning, and an adaptive attention mechanism, effectively mitigates cumulative error propagation that often occurs in multi-step forecasting tasks.

Figure 6 presents the MAE comparison across the five models for the 30-7 forecasting task. MAE serves as a direct measure of average prediction deviation and is a critical indicator of forecasting precision.

For the Palo Alto dataset, the proposed TF-EffBiGRU-AttNet records the lowest MAE of 0.0332, substantially outperforming Bi-LSTM (0.0937) and EfficientNet (0.14). This result confirms the model’s enhanced ability to capture detailed spatial-temporal features, reducing absolute deviations in long-term predictions. The Bi-GRU and GRU models, on the other hand, display significantly higher MAE values, revealing difficulties in managing error accumulation over multiple steps.

In the Perth dataset, although MAE values are slightly higher across all models due to more irregular demand fluctuations, the proposed model still achieves the smallest MAE of 0.0235, followed by Bi-LSTM (0.04) and EfficientNet (0.0968). The consistent superiority of TF-EffBiGRU-AttNet across both datasets demonstrates its strong predictive stability and generalization capability over extended forecasting windows.

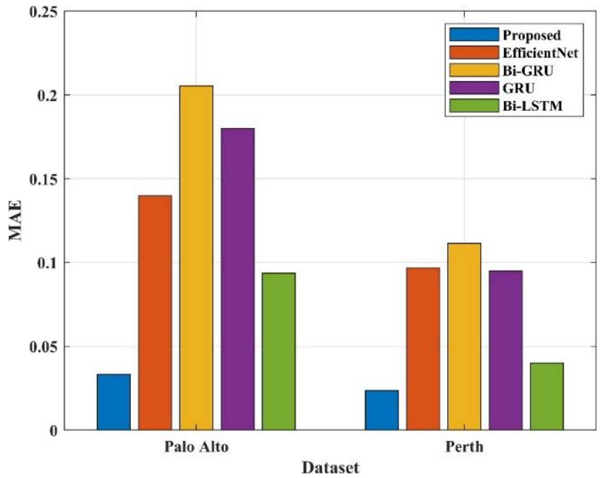

Figure 6. MAE performance comparison for Task 30-7.

Figure 7 illustrates the RMSE results for the five models under the same forecasting task. RMSE is sensitive to large prediction errors and reflects the model’s ability to suppress significant deviations.

The proposed model once again achieves the lowest RMSE values, recording 0.1654 for Palo Alto and 0.0824 for Perth. These values are markedly lower than those of Bi-LSTM (0.4861 and 0.2697) and much smaller compared to traditional GRU and Bi-GRU models. The significant improvement in RMSE demonstrates the model’s ability to maintain stable learning across longer prediction sequences, minimizing the amplification of forecast deviations.

For both datasets, the gap between TF-EffBiGRU-AttNet and the other models widens as the forecasting horizon increases. This indicates that the proposed hybrid approach, which combines efficient spatial encoding with bidirectional temporal learning, provides superior long-term stability compared to purely recurrent or convolutional baselines.

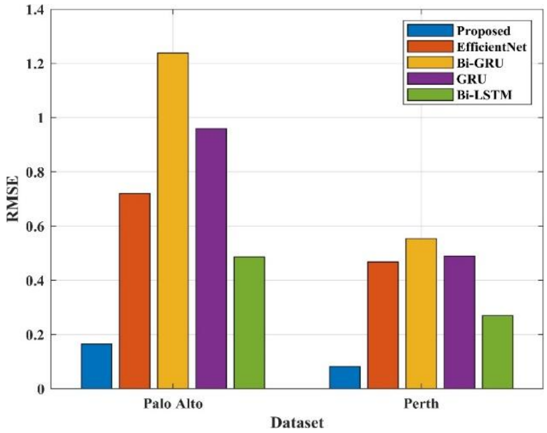

Figure 7. RMSE performance comparison for Task 30-7.

Figure 8 depicts the MAPE comparison among the five models. MAPE evaluates the forecasting performance in percentage terms and allows direct interpretability of prediction accuracy.

For both datasets, TF-EffBiGRU-AttNet yields the lowest MAPE—0.20% for Palo Alto and 0.37% for Perth—showing that the proposed model produces highly reliable predictions with minimal relative error. The improvement over EfficientNet (0.84% and 1.64%) and Bi-GRU (1.45% and 2.32%) highlights the benefit of combining adaptive attention with sequential temporal modeling.

The slightly higher MAPE for Perth again reflects greater demand variability and data sparsity. However, the margin between the proposed model and the next-best performer remains substantial, confirming the model's consistent ability to maintain accuracy under dynamic conditions.

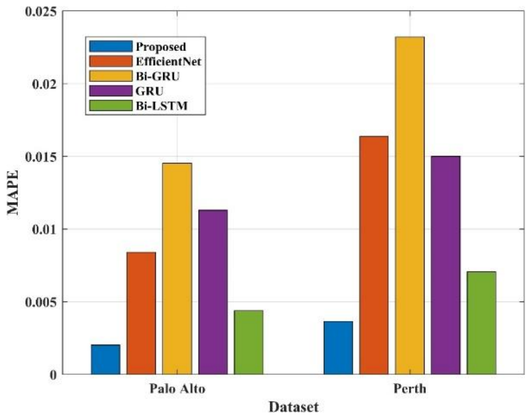

Figure 8. MAPE comparison for Task 30-7.

The results from table 3 and figures 6–8 clearly demonstrate that the proposed TF-EffBiGRU-AttNet model provides significant improvements in forecasting accuracy and error stability compared to traditional and hybrid baselines. The combination of dual-scale spatial encoding through EfficientNetV2 and bidirectional temporal modeling with attention enables the model to effectively learn both short- and long-term dependencies while suppressing cumulative prediction errors.

The superiority of TF-EffBiGRU-AttNet in the multi-step forecasting scenario highlights its ability to generalize across time horizons and complex spatial structures. The consistently lower error values in both datasets also confirm the role of the Self-Adaptive Hippopotamus Optimization Algorithm, which enhances model parameter tuning and convergence efficiency.

Overall, the findings from this task indicate that the proposed model not only achieves excellent short-term prediction accuracy but also maintains robust long-term forecasting stability. This makes it highly suitable for real-world applications in electric vehicle charging demand management, where continuous and accurate multi-step prediction is crucial for grid optimization and sustainable energy planning.

4.3. Optimization Performance Analysis

To enhance the learning efficiency and stability of the proposed TF-EffBiGRU-AttNet model, a comparative analysis of the Self-Adaptive Hippopotamus Optimization Algorithm (SA-HOA) was conducted against three well-known metaheuristic optimization algorithms: Particle Swarm Optimization (PSO), Grey Wolf Optimization (GWO), and the standard Hippopotamus Optimization Algorithm (HOA). This comparison evaluates the capability of each algorithm to optimize the model parameters effectively, reduce computational overhead, and achieve faster convergence during the training process.

Table 4 presents the comparative performance results based on six evaluation metrics: best fitness value (minimization), mean fitness value, standard deviation of fitness, convergence speed (iterations), computational time, and exploration–exploitation rates.

Table 4. Comparison of optimization algorithm performance.

Metric	PSO	GWO	HOA	SA-HOA (Proposed)
Best Fitness Value (Minimization)	0.0012	0.0009	0.0006	0.0003
Mean Fitness Value	0.0025	0.0018	0.0011	0.0005
Standard Deviation of Fitness	0.0008	0.0006	0.0004	0.0002
Convergence Speed (Iterations)	1200	1500	900	600
Computational Time (sec)	2.5	3.1	1.8	1.2
Exploration Rate (%)	65	75	80	90
Exploitation Rate (%)	60	70	85	95

The best fitness value represents the algorithm's ability to locate the global minimum of the objective function. The SA-HOA achieves the best result, with a minimum fitness value of 0.0003, outperforming HOA (0.0006), GWO (0.0009), and PSO (0.0012). This result highlights SA-HOA's superior global search capability, enabling it to identify optimal parameters that minimize forecasting errors in the TF-EffBiGRU-AttNet model.

The mean fitness value reflects the overall optimization consistency across multiple iterations. SA-HOA achieves the lowest mean fitness value (0.0005), demonstrating a stable and reliable convergence behavior. In contrast, PSO and GWO exhibit higher mean fitness values, indicating that their optimization paths fluctuate more during the search process. This consistency is crucial for deep learning models, where stable parameter convergence directly affects training efficiency and prediction accuracy.

The standard deviation of fitness measures the variability of fitness values during the optimization process. SA-HOA again produces the smallest deviation (0.0002), showing a more uniform convergence trajectory and reduced oscillation around the optimal solution. This low variance indicates that the algorithm effectively avoids premature convergence and maintains a steady optimization process, even in high-dimensional search spaces.

The convergence speed and computational time metrics assess how quickly and efficiently the algorithms reach optimal solutions. SA-HOA converges in 600 iterations, significantly faster than HOA (900 iterations), PSO (1200 iterations), and GWO (1500 iterations). The improved convergence speed results from the algorithm's self-adaptive strategy, which dynamically adjusts the balance between exploration and exploitation based on the current search stage.

In terms of computational time, SA-HOA completes optimization in 1.2 seconds, which is the shortest among all tested algorithms. The reduction in time consumption demonstrates the efficiency of the adaptive control mechanism that eliminates redundant searches and focuses computation on promising solution regions. This characteristic makes SA-

HOA particularly well-suited for large-scale or real-time forecasting systems, where computational efficiency is essential.

The exploration rate quantifies the algorithm's ability to explore new regions in the search space, while the exploitation rate indicates its ability to refine known promising regions. Achieving an optimal balance between the two is critical for avoiding local minima and ensuring global optimality.

SA-HOA achieves the highest exploration rate of 90% and the highest exploitation rate of 95%, outperforming HOA (80% and 85%), GWO (75% and 70%), and PSO (65% and 60%). This improvement is attributed to SA-HOA's self-adaptive control mechanism, which continuously adjusts the search parameters—such as movement factors and local influence coefficients—based on iteration progress and current population diversity. The algorithm starts with a high exploration emphasis to cover the global search space and gradually increases exploitation to fine-tune the solution near convergence.

The results from table 4 clearly demonstrate that the Self-Adaptive Hippopotamus Optimization Algorithm provides significant improvements over traditional metaheuristic algorithms across all evaluation metrics. Its adaptive nature enables it to dynamically respond to changes in the optimization landscape, leading to faster convergence, reduced computational effort, and more stable performance.

These characteristics directly translate to enhanced model performance in the TF-EffBiGRU-AttNet framework. The optimized parameters achieved by SA-HOA improve both prediction accuracy and generalization, while reducing overfitting and training time. Moreover, the algorithm's strong exploration–exploitation balance ensures that the model can effectively handle the nonlinear and dynamic nature of spatio-temporal EV charging demand.

Overall, the integration of SA-HOA not only enhances the computational efficiency of the proposed framework but also strengthens its robustness and adaptability for complex, real-world forecasting scenarios in smart energy management systems.

4.4. Discussion

The experimental results confirm that the integration of Two-Fold EfficientNetV2, Bidirectional GRU (BiGRU), Attention Mechanism, and Self-Adaptive Hippopotamus Optimization Algorithm (SA-HOA) significantly enhances forecasting performance compared to traditional deep learning models [16], [18], [25]. The proposed TF-EffBiGRU-AttNet model effectively captures both spatial and temporal dependencies in electric vehicle (EV) charging demand, resulting in superior predictive accuracy and generalization capability across multiple datasets [19], [20], [23].

The Two-Fold EfficientNetV2 module strengthens spatial feature extraction by employing compound scaling on network depth, width, and resolution [17]. This enables the model to learn both local and global spatial correlations efficiently. Consequently, the model can represent complex spatial relationships influenced by station distribution, urban topology, and traffic density, which directly affect energy consumption behavior [23]. When integrated with the BiGRU, the model captures bidirectional temporal dependencies, allowing it to model sequential patterns across different time horizons such as hourly, daily, and seasonal variations [19], [21]. This capability is crucial for understanding temporal fluctuations in EV charging activities driven by user behavior and environmental factors [24].

The Attention Mechanism further improves interpretability and accuracy by adaptively assigning higher weights to spatial and temporal features that have stronger impacts on prediction outcomes [25]. This selective focus allows the model to identify critical time periods and station clusters that contribute most significantly to demand fluctuations. At the same time, the SA-HOA optimization algorithm enhances learning efficiency by automatically tuning model parameters during training. Its self-adaptive behavior ensures faster convergence, reduced overfitting, and improved balance between exploration and exploitation [20], [22].

Across all evaluation metrics, including MAE, RMSE, and MAPE, the TF-EffBiGRU-AttNet consistently outperforms baseline models such as EfficientNet, Bi-GRU, GRU, and Bi-LSTM [16], [18], [21]. The lower error rates in the Palo Alto dataset demonstrate the model's strong capability in capturing structured and stable energy demand patterns. The slightly higher errors in the Perth dataset are attributed to greater spatial heterogeneity, environmental variability, and user behavior uncertainty [23], [24]. Despite these challenges, the TF-EffBiGRU-AttNet maintains superior accuracy and stability, indicating its ability to generalize effectively across diverse geographic and temporal contexts.

Overall, the proposed TF-EffBiGRU-AttNet framework demonstrates outstanding performance in both short-term and long-term forecasting tasks. Its hybrid structure achieves a balance between accuracy, interpretability, and computational efficiency, making it highly suitable for deployment in real-world electric vehicle charging infrastructures [19], [25]. The model's predictive capability supports proactive energy management, efficient charging station allocation, and optimized grid load balancing [17], [21], [22]. These outcomes validate the potential of TF-EffBiGRU-AttNet as a robust and scalable framework for intelligent energy demand forecasting in EV charging networks, contributing to the advancement of smart grid systems and sustainable urban transportation [23], [26].

5. Conclusion

The experimental findings confirm that the proposed TF-EffBiGRU-AttNet model achieves superior performance in forecasting electric vehicle (EV) charging demand compared to baseline models, as demonstrated using the Palo Alto and Perth datasets [16], [19], [23]. The model consistently delivers more accurate predictions with lower errors across all evaluation metrics, including MAE, RMSE, and MAPE. This improvement is attributed to the integration of the Two-Fold EfficientNetV2, Bidirectional GRU (BiGRU), and Attention Mechanism, which together enable the model to effectively capture complex spatial and temporal dependencies [17], [21], [25].

The Two-Fold EfficientNetV2 module extracts multi-scale spatial features, allowing the model to represent variations in EV charging behavior across different geographical areas [17], [23]. The BiGRU component captures dynamic temporal relationships in both forward and backward directions, effectively modeling daily and seasonal consumption trends [19], [21]. The Attention Mechanism enhances the model's interpretability and accuracy by prioritizing the most relevant features in space and time [25]. Furthermore, the integration of the Self-Adaptive Hippopotamus Optimization Algorithm (SA-HOA) ensures efficient parameter tuning and stable convergence, reducing overfitting and enhancing the model's adaptability to heterogeneous datasets [20], [22].

The results also reveal that longer forecasting horizons, such as multi-step prediction, lead to higher cumulative errors for all models, which is a common limitation in temporal prediction tasks [18], [24]. Despite this challenge, the TF-EffBiGRU-AttNet demonstrates resilience and robustness, maintaining lower error rates than competing models even under complex forecasting conditions. The sensitivity of the model to hyperparameters, particularly the learning rate, emphasizes the importance of optimized configuration during training to achieve maximum performance.

In practical applications, the TF-EffBiGRU-AttNet model can serve as an intelligent tool for strategic placement of EV charging stations, smart grid integration, and optimized power distribution management [21], [23], [26]. Its ability to model spatial, seasonal, and nonlinear energy consumption patterns supports data-driven decision-making for sustainable energy systems. However, potential challenges remain in improving real-time adaptability, generalization to unseen environments, and data privacy protection when deployed in large-scale, real-world networks. Addressing these aspects through adaptive learning and decentralized optimization approaches could further enhance the model's applicability and impact in future research.

Overall, the proposed TF-EffBiGRU-AttNet framework provides a robust, accurate, and scalable solution for spatio-temporal energy demand forecasting in electric vehicle charging networks. Its integration of advanced deep learning and optimization methods offers a promising direction for the development of intelligent and sustainable energy management systems in modern smart cities [17], [19], [25].

6. Declarations

6.1. Author Contributions

Author Contributions: Conceptualization, S.P., S.A.M., and G.S.; Methodology, S.P. and M.B.; Software, G.S. and S.A.M.; Validation, S.A.M. and M.B.; Formal Analysis, S.P.; Investigation, G.S. and S.A.M.; Resources, S.A.M. and M.B.; Data Curation, G.S.; Writing—Original Draft Preparation, S.P.; Writing—Review and Editing, M.B. and S.A.M.; Visualization, S.A.M. All authors have read and agreed to the published version of the manuscript.

6.2. Data Availability Statement

The data presented in this study are available on request from the corresponding author.

6.3. Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

6.4. Institutional Review Board Statement

Not applicable.

6.5. Informed Consent Statement

Not applicable.

6.6. Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Reference

[1] A. V. Sreekumar and R. R. Lekshmi, "Electric vehicle charging station demand prediction model deploying data slotting," *Results in Engineering*, vol. 24, no. Dec., pp. 1–12, 2024.

[2] Y. Li, R. Xie, C. Li, Y. Wang, and Z. Dong, "Federated Graph Learning for EV Charging Demand Forecasting with Personalization Against Cyberattacks," *arXiv (Cornell University)*, vol. 2024, no. Apr., pp. 1–18, 2024.

[3] N. Ji, R. Zhu, Z. Huang, and L. You, "An urban-scale spatiotemporal optimization of rooftop photovoltaic charging of electric vehicles," *Urban Informatics*, vol. 3, no. 1, pp. 1–15, 2024.

[4] H. Liu, Z. Xing, Q. Zhao, Y. Liu, and P. Zhang, "An Orderly Charging and Discharging Strategy of Electric Vehicles Based on Space–Time Distributed Load Forecasting," *Energies*, vol. 17, no. 17, pp. 4284–4298, 2024.

[5] C. Yao, S. Chen, M. Salazar, and Z. Yang, "Joint Routing and Charging Problem of Electric Vehicles With Incentive-Aware Customers Considering Spatio-Temporal Charging Prices," *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 11, pp. 12215–12226, 2023.

[6] J. Li, S. Tian, N. Zhang, G. Liu, Z. Wu, and W. Li, "Optimization Strategy for Electric Vehicle Routing under Traffic Impedance Guidance," *Applied Sciences*, vol. 13, no. 20, pp. 11474–11486, 2023.

[7] J. Feng, X. Chang, Y. Fan, and W. Luo, "Electric Vehicle Charging Load Prediction Model Considering Traffic Conditions and Temperature," *Processes*, vol. 11, no. 8, pp. 2256–2269, 2023.

[8] M. W. Mudiyanselage, "A Multi-Agent Framework for Electric Vehicles Charging Power Forecast and Smart Planning of Urban Parking Lots," *IEEE Transactions on Transportation Electrification*, vol. 2023, no. Jan., pp. 1–14, 2023.

[9] R. Luo, Y. Song, L. Huang, Y. Zhang, and R. Su, "AST-GIN: Attribute-Augmented Spatiotemporal Graph Informer Network for Electric Vehicle Charging Station Availability Forecasting," *Sensors*, vol. 23, no. 4, pp. 1975–1989, 2023.

[10] Y. Wang, D. Zhao, Y. Ren, D. Zhang, and H. Ma, "SPAP: Simultaneous Demand Prediction and Planning for Electric Vehicle Chargers in a New City," *ACM Transactions on Knowledge Discovery from Data*, vol. 17, no. 4, pp. 1–25, 2023.

[11] Y. N. Wang, Z. J. Zhang, A. Ping, R. J. Wang, and D. Q. Gong, "Optimizing electric vehicle charging strategies using multi-layer perception-based spatio-temporal prediction," *Advances in Production Engineering & Management*, vol. 19, no. 4, pp. 443–459, 2024.

[12] T. van Etten, V. Degeler, and D. Luo, "Large-Scale Forecasting of Electric Vehicle Charging Demand Using Global Time Series Modeling," *Proceedings of the 7th International Conference on Vehicle Technology and Intelligent Transport Systems*, vol. 2024, no. Jan., pp. 40–51, 2024.

[13] L. You, "FMGCN: Federated Meta Learning-Augmented Graph Convolutional Network for EV Charging Demand Forecasting," *IEEE Internet of Things Journal*, vol. 2024, no. Jan., pp. 1–13, 2024.

[14] P. Chen, J. Qin, J. Dong, L. Ling, X. Lin, and H. Ding, "Electric vehicle charging demand forecasting at charging stations under climate influence," *IET Power Electronics*, vol. 18, no. 1, pp. 1–14, 2025.

[15] R. Tian, J. Wang, Z. Sun, J. Wu, X. Lu, and L. Chang, “Multi-Scale Spatial-Temporal Graph Attention Network for Charging Station Load Prediction,” *IEEE Access*, vol. 13, no. 2025, pp. 29000–29017, 2025.

[16] S. Wang, Y. Li, C. Shao, P. Wang, A. Wang, and C. Zhuge, “An adaptive spatio-temporal graph recurrent network for short-term electric vehicle charging demand prediction,” *Applied Energy*, vol. 383, no. Apr., pp. 1–15, 2025.

[17] H. Kuang, K. Deng, L. You, and J. Li, “Citywide electric vehicle charging demand prediction approach considering urban region and dynamic influences,” *Energy*, vol. 320, no. Apr., pp. 1–14, 2025.

[18] Y. Chen, Z. Tang, Y. Cui, W. Rao, and Y. Li, “Electric Vehicle Charging Demand Prediction Model Based on Spatiotemporal Attention Mechanism,” *Energies*, vol. 18, no. 3, pp. 687–701, 2025.

[19] X. Yang, “A spatiotemporal distribution prediction model for electric vehicles charging load in transportation power coupled network,” *Scientific Reports*, vol. 15, no. 1, pp. 1–14, 2025.

[20] R. Marlin, R. Jurdak, and A. Abuadbba, “H-FLTN: A Privacy-Preserving Hierarchical Framework for Electric Vehicle Spatio-Temporal Charge Prediction,” *arXiv.org*, vol. 2025, no. Feb., pp. 1–16, 2025.

[21] Q. Han and X. Li, “A Vertical Federated Learning Method for Electric Vehicle Charging Station Load Prediction,” *Processes*, vol. 13, no. 2, pp. 468–482, 2025.

[22] Y. Zhuang, L. Cheng, N. Qi, X. Wang, and Y. Chen, “Real-time hosting capacity assessment for electric vehicles,” *Applied Energy*, vol. 380, no. Dec., pp. 1–14, 2024.

[23] S. Hou, X. Zhang, and H. Yu, “Electric Vehicle Charging Load Prediction Considering Spatio-Temporal Node Importance Information,” *Energies*, vol. 17, no. 19, pp. 4840–4854, 2024.

[24] X. Hu, “GCN-Transformer-Based Spatio-Temporal Load Forecasting for EV Battery Swapping Stations under Differential Couplings,” *Electronics*, vol. 13, no. 17, pp. 3401–3415, 2024.

[25] Y. Chen, M. Wang, Y. Wei, X. Huang, and S. Gao, “Multi-Encoder Spatio-Temporal Feature Fusion Network for Electric Vehicle Charging Load Prediction,” *Journal of Intelligent & Robotic Systems*, vol. 110, no. 3, pp. 1–16, 2024.

[26] M. S. Hasibuan, R. Z. A. Aziz, D. A. Dewi, T. B. Kurniawan, and N. A. Syafira, “Recommendation Model for Learning Material Using the Felder Silverman Learning Style Approach,” *HighTech and Innovation Journal*, vol. 4, no. 4, pp. 811–820, 2023.