Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Context-Enriched Graph Neural Networks for EV Charging Station Energy Demand Forecasting

Douaidi Lydia

douaidilydia98@gmail.com

Université Bourgogne Europe https://orcid.org/0009-0008-7606-930X

Senouci Sidi-Mohammed

Université Bourgogne Europe

Zekri latifa-Aya

École supérieure en informatique de Sidi Bel Abbès

Fouzi Harrou

King Abdullah University of Science and Technology (KAUST)

Research Article

Keywords: Charging Stations, Electric Vehicles, Energy Forecasting, Artificial Intelligence, Graph Neural Network

Posted Date: June 24th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10128538/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Context-Enriched Graph Neural Networks for EV Charging Station Energy Demand Forecasting

Lydia Douaidi[a], Latifa Aya Zekri[a,b], Sidi-Mohammed Senouci[a], Fouzi Harrou[c]

[a] *Université Bourgogne Europe, DRIVE Lab, Nevers, France*
E-mail: {firstname.lastname}@ube.fr
[b] *École Supérieure en Informatique 08 Mai 1945, Sidi Bel-Abbès, Algeria*
E-mail: la.zekri@esi-sba.dz
[c] *King Abdullah University of Science and Technology (KAUST), Computer, Electrical and Mathematical Sciences and Engineering (CEMSE) Division, Thuwal 23955-6900, Saudi Arabia; E-mail: fouzi.harrou.2016@ieee.org*

Abstract

Efficient management of electric vehicle charging station (EVCS) networks is essential to ensure grid stability and sustainable urban mobility. Accurate load forecasting supports operators in balancing energy consumption, handling variable demand, and planning infrastructure expansion. However, existing forecasting methods often neglect spatial dependencies among stations, or insufficiently exploit heterogeneous urban data, limiting prediction accuracy. This paper proposes a spatio-temporal graph neural network (STGNN) framework that integrates heterogeneous features, historical load profiles, points-of-interest distributions, and population density into static spatial graphs for EVCS forecasting. Four advanced STGNN architectures, Multivariate Time Graph Neural Network (MTGNN), Attention-Based Spatio-Temporal Graph Convolutional Network (ASTGCN), Spatio-Temporal Convolution (STConv), and Attention Temporal Graph Convolutional Network (A3TGCN), are systematically benchmarked against temporal baselines, Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU). Experiments on two geographically diverse datasets, Palo Alto, USA (urban), and Yonne, France (semi-rural), demonstrate that the best STGNN reduces test root mean squared error to 8.32 and 15.16, compared with 10.65/18.24 for LSTM and 11.31/20.36 for GRU, with R^2 up to 0.86. Graph-based spatio-temporal modeling enables accurate, efficient, and generalizable EVCS forecasting across diverse urban settings.

Keywords: Charging Stations; Electric Vehicles; Energy Forecasting; Artificial Intelligence; Graph Neural Network.

Nomenclature

Abbreviations	
EV	Electric Vehicle
EVCS	Electric Vehicle Charging Station
CS	Charging Station
CSO	Charging Station Operator
POI	Point of Interest
DL	Deep Learning
ML	Machine Learning
CNN	Convolutional Neural Network
RNN	Recurrent Neural Network
LSTM	Long Short-Term Memory network
GRU	Gated Recurrent Unit
GNN	Graph Neural Network
GCN	Graph Convolutional Network
STGNN	Spatio-Temporal Graph Neural Network
MTGNN	Multivariate Time Graph Neural Network
ASTGCN	Attention-based Spatio-Temporal GNN
STConv	Spatio-Temporal Convolutional Network
A3TGCN	Attention Temporal Graph Convolutional Network
STMGCN	Spatiotemporal Multi-Graph Convolutional Network
VSTF-GNN	Variational Spatiotemporal Factorized GNN
GAN	Generative Adversarial Network
WGAN-GP	Wasserstein GAN with Gradient Penalty
Metrics	
MSE	Mean Squared Error
RMSE	Root Mean Squared Error
MAE	Mean Absolute Error
R^2	Coefficient of Determination
Datasets	
Palo Alto	Urban dataset collected from Palo Alto, USA
Yonne (SDEY)	Semi-rural dataset collected from Yonne, France

1. Introduction

The rapid growth of electric mobility, driven by CO_2 reduction goals, has intensified demand for charging infrastructure and placed heavy burdens on power grids [1]. The widespread deployment of electric vehicle charging stations (EVCSs) has led to significant increases in power consumption, particularly during peak periods, which could cause losses, grid overload, and reduced power quality [2]. For instance, in 2024, the global electric vehicle fleet consumed nearly 180 TWh of electricity, marking an increase of almost 60% compared to the previous year [3]. These trends highlight the need for accurate energy demand forecasting to mitigate grid stress and support sustainable EVCS expansion.

Accurate load forecasting has become indispensable to manage rising electricity demand amid growing electric mobility [4]. It optimizes grid operations and reinforces stability, enabling sustainable electric vehicle (EV) integration while preserving efficiency and power quality. As renewable sources, such as solar and wind, expand, precise forecasting helps Charging Station Operators (CSOs) balance supply and demand, manage generation variability, and ensure reliable energy availability [5, 6]. Moreover, by incorporating live data, traffic, weather, energy prices, and real-time grid load, EV charging behaviors can be better predicted, improving scheduling and reducing range anxiety [7, 8]. Overall, accurate forecasting enhances short-term operations and long-term grid planning, forming a foundation for resilient and sustainable energy systems in the era of electrified transportation [9].

The efficient operation of EV charging networks requires sophisticated management strategies. Over the last decade, researchers have explored various approaches, including statistical and mathematical models [10], traditional optimization techniques such as linear programming and queuing theory, and ML techniques, to address challenges like demand forecasting [11], infrastructure optimization, charging occupancy prediction [12], and station recommendations [1]. Within this range of methodological approaches, accurate demand forecasting emerges as a key component, as future charging usage patterns exhibit strong temporal correlations with historical demand [13, 14]. Traditional time series approaches, Moving Average (MA), Autoregressive (AR), and Autoregressive Integrated Moving Average (ARIMA) [15], have been widely used for their interpretability and ability to capture seasonal or cyclic patterns [16]. However, they are limited in representing the complex

spatio-temporal dependencies of EV charging demand. To overcome these shortcomings, hybrid approaches have emerged, such as ARIMA–ML models that combine statistical interpretability with nonlinear learning [17], multi-scale spatio-temporal frameworks that better capture cross-station dependencies [18], and probabilistic hierarchical models for coherent multi-station demand prediction [19]. More recently, Deep Learning (DL) and big data have enabled advanced methods capable of modeling intricate nonlinear dynamics and contextual factors in EV charging demand. Convolutional and recurrent neural network models have demonstrated notable improvements in forecasting accuracy [20, 21], while attention-based and transformer-enhanced architectures further improved robustness across diverse charging scenarios [22]. Despite recent advances, accurate and reliable forecasting still requires frameworks that can jointly capture spatial correlations and temporal dynamics.

This paper addresses the pressing challenge of accurately forecasting energy consumption in EVCS networks. Unlike approaches centered on minute-level control, our study emphasizes reproducibility, adaptability across heterogeneous urban contexts, and reduced data requirements, thereby enhancing its practical applicability for infrastructure planning and grid coordination. We introduce a tailored spatio-temporal Graph Neural Network (GNN) framework that explicitly models both the temporal evolution of charging demand and the spatial interdependencies between stations. Our approach proposes multiple GNN architectures to capture complex energy patterns, while enriching each station's representation with contextual urban features to better reflect real-world charging behaviors. They key contributions of this research are as follows.

- A station-level, day-ahead forecasting framework is proposed using advanced spatio-temporal Graph Neural Network (STGNN) architectures. The study investigates four representative models: Multivariate Time Graph Neural Network (MTGNN), which integrates temporal convolutions with graph structure learning; Attention-Based Spatio-Temporal Graph Convolutional Network (ASTGCN), which applies spatial and temporal attention to capture dynamic dependencies; Spatio-Temporal Convolutional Network (STConv), which efficiently models local temporal patterns with graph convolutions; and Attention Temporal Graph Convolutional Network (A3TGCN), which combines gated temporal units with attention mechanisms for sequence learning. These architectures are selected to represent complementary families

of STGNNs, enabling a comprehensive evaluation of spatio-temporal learning paradigms for EVCS energy demand forecasting.

- The proposed framework integrates multi-source data, including spatial relationships, historical usage patterns, and contextual urban features such as Points-of-Interest (POI) distributions and population density. Temporal inputs are represented by sliding windows of charging demand, while spatial dependencies are encoded in a weighted adjacency matrix that reflects inter-station connectivity. Node feature vectors combining these heterogeneous data sources are then processed through STGNN models with stacked graph convolution layers, which enhance generalization, accuracy, and stability. This comprehensive data fusion captures demand fluctuations shaped by spatial connectivity, user behavior, and demographic factors, thereby improving forecasting accuracy and strengthening the model's practical relevance for EVCS management.

- Two real-world datasets, Palo Alto, USA (urban) and Yonne, France (semi-rural), are used to evaluate the proposed framework, ensuring assessment across contrasting geographic and demographic contexts. Performance is compared against established temporal baselines, Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU), with multiple complementary metrics providing a fair and comprehensive evaluation of forecasting accuracy and generalizability.

The remainder of this paper is organized as follows. Section II reviews related work. Section III presents the proposed framework. Section IV describes the dataset. Section V reports experiments and evaluates performance. Section VI concludes the study and outlines future research directions.

2. Related Work

With the rapid growth of electric vehicle adoption, research has increasingly focused on optimizing charging infrastructure to meet rising energy demands. The literature presents a variety of approaches, including theoretical modeling, simulations, and data-driven techniques, to forecast load dynamics and manage demand at charging stations (CSs) [23, 24, 14]. Statistical analyses have been applied to estimate peak usage periods [23]. More

recent data-driven models use historical charging patterns to forecast future demand and support grid load management [24, 25]. This section reviews EVCS energy forecasting methods and highlights their role in improving infrastructure resilience, operational efficiency, and the integration of EVs into sustainable urban energy systems.

Accurately predicting energy demand at EV charging stations is crucial for optimizing resource allocation and infrastructure planning. Statistical time-series models such as ARIMA and Seasonal ARIMA (SARIMA) remain widely used for their simplicity and interpretability [15, 26]. However, they are limited in capturing nonlinear dynamics and complex spatio-temporal dependencies in charging demand. More recently, DL has emerged as a powerful tool for time-series forecasting. For instance, a Bidirectional Long Short-Term Memory (Bi-LSTM) model was applied for load prediction in [27], while Convolutional Neural Networks (CNNs) have been deployed on grid-based maps to capture spatial correlation patterns [28]. A hybrid transfer learning framework combining CNN and Bi-LSTM was further proposed in [29], improving forecasting accuracy under limited data conditions. These advances highlight the role of hybrid models in mitigating data scarcity while maintaining reliable prediction performance.

However, conventional DL methods, such as CNNs and Recurrent Neural Networks (RNNs), often struggle to capture the heterogeneous spatial correlations inherent in EV charging demand (e.g., variations in usage by location, POI, urban layout) [25, 30]. In contrast, graph-based representations and Graph Convolutional Networks (GCNs) are increasingly adopted for their ability to model complex spatial dependencies and improve forecasting accuracy [14, 20]. In a comprehensive survey, Rashid et al. [9] reviewed statistical and ML-based demand forecasting methods, highlighting the growing role of machine learning in enhancing prediction accuracy and operational efficiency. For example, Hüttel et al. [21] proposed a DL–GNN model to forecast daily energy demand at public CSs using data from Palo Alto, USA, while Wang et al. [20] developed a spatio-temporal GCN that achieved a 12.69% accuracy improvement. A more advanced framework, the Variational Spatiotemporal Factorized GNN (VSTF-GNN), was introduced in [31] to jointly address EV usage prediction and pattern recognition at the network level, integrating graph-based spatial modeling with time-series forecasting and variational Bayesian methods to handle missing data. Generative Adversarial Networks (GANs) have also been employed in this

domain: a Wasserstein GAN with Gradient Penalty (WGAN-GP) was used to generate highly variable spatio-temporal charging load scenarios, supporting robust training and evaluation [32], while [33] applied GANs in a 3D modeling framework to capture complex EV demand patterns.

Further advances combine GCNs with other DL models. A hybrid GCN–BiLSTM model was designed to capture both static and dynamic attributes of EVCSs [34], while Fahim et al. [35] integrated GCN with LSTM to jointly exploit spatial and temporal features for demand forecasting. More recently, the Spatiotemporal Multi-Graph Convolutional Network (STMGCN) [36] has incorporated temporal signals and spatial data from neighboring stations, achieving state-of-the-art accuracy on real-world datasets.

Despite these advances, most existing models still overlook critical contextual data such as POIs and demographic factors, limiting their predictive accuracy. In addition, much of the literature emphasizes aggregate demand forecasting over time intervals (e.g., hourly or daily averages) rather than delivering fine-grained, station-level predictions. These gaps underscore the need for approaches that capture both spatio-temporal dependencies and localized charging behaviors. GNNs, particularly when enhanced with attention mechanisms and enriched node features, provide a powerful framework to address these challenges, enabling more accurate forecasts and deeper insights into EVCS demand dynamics.

Table 1 further illustrates how diverse methodological families, including DL models, graph-based approaches, and generative models, offer complementary strengths but also face persistent challenges, emphasizing the necessity of integrated frameworks in EVCS energy demand forecasting.

3. Research Framework

This section presents the proposed framework, which exploits STGNNs for energy-demand forecasting.

State-of-the-art models capture spatial and temporal patterns in charging demand [28], but they often emphasize short-term forecasts at fine-grained intervals. While useful for real-time load balancing, such approaches offer limited value for long-term planning and infrastructure optimization. Instead, this study forecasts daily demand at the station level, directly supporting capacity planning, investment decisions, and grid coordination. To this end, STGNNs are employed to jointly model inter-station relationships

Table 1: Comparative analysis of selected methods for EVCS energy demand forecasting.

Method	**Advantages**	**Challenges**	**Key studies**
RNNs, CNNs	Capture temporal patterns and nonlinear relationships; handle large datasets and time series.	Computationally intensive; overlook spatial dependencies; risk of overfitting without regularization.	[27, 28, 29]
GNNs	Effectively model spatial dependencies; capture inter-station relationships (e.g., proximity).	Graph construction is non-trivial; limited scalability on large networks; sensitivity to graph sparsity.	[21, 20, 34, 9, 36]
GANs	Generate realistic scenarios for data augmentation; helpful for imbalance and simulating diverse conditions.	Computationally demanding; training instability; limited interpretability.	[33, 32]

and temporal dynamics [21]. EVCSs are represented as nodes in a spatial graph enriched with contextual features such as POIs, population density, and historical usage. This enables the model to capture both spatio-temporal correlations and socio-demographic drivers of charging demand.

The following subsection outlines the principles of STGNNs as the basis of the approach. The construction of the static spatial graph is then detailed, where nodes represent charging stations and edges capture their interrelationships. A sliding-window mechanism models the temporal evolution of demand, followed by the presentation of advanced STGNN architectures for daily energy demand forecasting.

3.1. Graph Construction and Spatio-Temporal Representation

STGNNs have emerged as a powerful tool for capturing complex spatio-temporal dependencies in various real-world applications. By extending traditional GNNs to incorporate temporal dynamics, they enable the joint analysis of spatial relationships and time-evolving patterns within a unified framework [37, 38, 39]. This capability enables STGNNs to model intricate node

interactions over time, proving highly effective in domains such as traffic prediction [40], environmental monitoring [41], and, more recently, electric vehicle charging infrastructure management [37], where both temporal patterns and spatial dependencies are critical for accurate forecasting and optimization.

The operation of EVCSs is inherently interdependent: demand at one station is often influenced by nearby stations, surrounding infrastructure, and demographic factors [42]. For instance, a surge in charging at a station near a shopping center or workplace may reduce demand at adjacent stations. Such spatial dependencies are essential for accurate forecasting, as they reflect how charging behavior propagates across the urban environment.

To capture these interactions, a static spatial graph is constructed in which each node represents a charging station (CS) and edges encode their relationships. Edges are defined using road-network connectivity (e.g., travel time between stations) to account for physical proximity, while functional proximity reflects correlations in demand across stations serving similar urban roles, even when they are not geographically adjacent. In a static spatio-temporal graph, the topology (i.e., nodes and edges) remains fixed over time, whereas the node states evolve dynamically to represent temporal variations in charging demand [43]. Figure 1 illustrates the integration of spatial and temporal relationships within this framework. The edges are directed, forming a network that represents spatial flows and directional influences, such as load redistribution when a station is overloaded or temporarily offline. This configuration allows the model to capture both spatial dependencies and directional dynamics, which are crucial for accurate demand forecasting and effective network balancing.

Within this graph, STGNN models employ message passing to propagate information between connected stations [44]. This mechanism allows each station to iteratively exchange information with its neighbors, enabling the models to capture how local demand fluctuations may influence surrounding nodes. Formally, the process consists of two steps:

- **Message Function:** This function computes a message from one node to another as a function of the features of the nodes and the attributes of the connecting edge, represented as:

$$m_{i \to j} = M(h_i, h_j, e_{ij}) \tag{1}$$

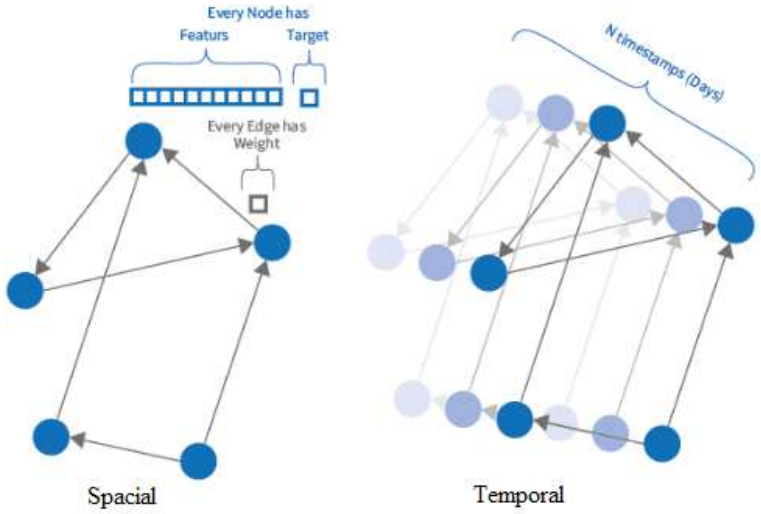


Figure 1: Node representations with their associated features within the spatio-temporal graph. Left: spatial graph where nodes carry features and edges encode weighted inter-site relations; Right: temporal unrolling into N daily snapshots linking node copies across time to model spatial dependencies and temporal dynamics.

- **Update Function:** This function updates the state of each node based on the messages it has received, represented as:

$$h_i^{new} = U\left(h_i, \sum_j m_{j\to i}\right) \tag{2}$$

  Here, h_i and e_{ij} are the feature vectors of node i and edge ij, respectively, and $m_{j\to i}$ is the message from node j to node i.

Each node in the graph, representing a CS, has a state that evolves over time, effectively modeling the station's energy demand at sequential time intervals and creating a dynamic time series. Temporal information is incorporated like traditional time series forecasting methods, capturing patterns and fluctuations in demand over time [45]. The evolution of each node's state is modeled such that its future state is predicted as a function of its historical states, formalized by:

$$h_i^{t+1} = f(h_i^t, h_i^{t-1}, \ldots, h_i^1). \tag{3}$$

3.2. Forecasting Framework with Sliding-Window Mechanism

After constructing the spatio-temporal graph, we apply a sliding window mechanism to generate the input sequences for the forecasting model, as illustrated in Fig. 2. At each forecasting step t, the model receives a sequence of historical observations from the past w time steps, denoted as:

$$\mathbf{X}_{t-w+1:t} \in \mathbb{R}^{w \times N \times F}, \tag{4}$$

where w is the length of the historical window, N is the total number of nodes in the graph, and F is the number of features per node (e.g., energy demand, temporal indicators , etc.).

This representation encodes the temporal evolution of each node while preserving the spatial relationships among nodes within the charging network.

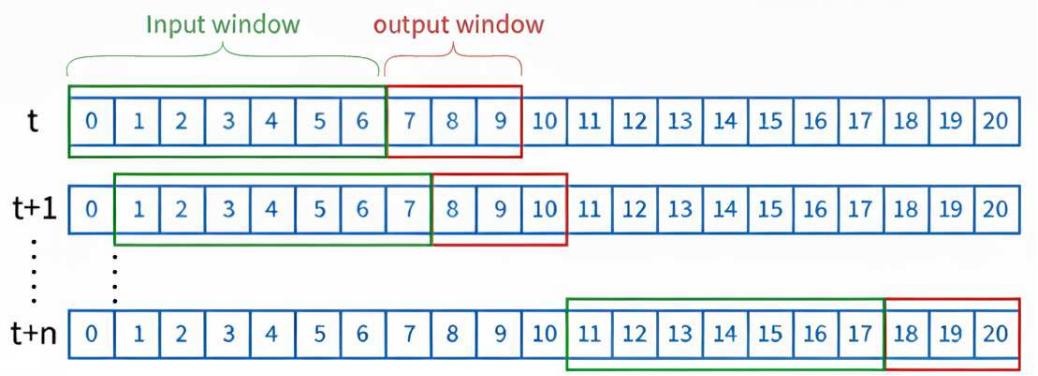


Figure 2: Sliding window mechanism for spatio-temporal forecasting of energy demand.

The model processes this spatio-temporal input to simultaneously capture temporal and spatial dependencies. The hidden state representation at time t, denoted as:

$$\mathbf{H}_t = g_{\mathrm{ST}}\left(\mathbf{X}_{t-w+1:t}; \Theta\right) \tag{5}$$

is generated by the spatio-temporal encoder $g_{\mathrm{ST}}(\cdot)$, parameterized by Θ, and encapsulates both temporal and spatial correlations across all stations.

The forecasting head then maps the learned representation $\mathbf{H}_t$ to the predicted future demand:

$$\hat{\mathbf{X}}_{t+1:t+w} = f(\mathbf{H}_t) \tag{6}$$

where:

$$\hat{\mathbf{X}}_{t+1:t+w} \in \mathbb{R}^{w \times N \times F} \tag{7}$$

represents the predicted energy consumption for all N stations over the next w time steps.

To emphasize the node-level temporal dynamics, the prediction for each station i can be expressed as:

$$\hat{\mathbf{x}}_i^{t+1:t+w} = f_i\left(\mathbf{x}_i^{t-w+1:t}, \{\mathbf{x}_j^{t-w+1:t} \mid j \in \mathcal{N}(i)\}\right) \tag{8}$$

where:

- $\mathbf{x}_i^{t-w+1:t}$ is the historical demand profile of station i,
- $\mathcal{N}(i)$ is the neighborhood of station i in the spatial graph, and
- $f_i(\cdot)$ captures how the temporal evolution of i is influenced by both its own history and the activity of nearby stations.

This formulation enables the model to learn complex spatio-temporal correlations, making it robust for forecasting network-wide demand and identifying dynamic interactions within the charging infrastructure.

3.3. Model Architectures and Baselines

To predict energy usage accurately at EVCSs within the proposed spatio-temporal graph framework, four advanced graph-based models are considered. Each model applies a distinct approach to processing spatial and temporal dependencies, enabling the assessment of different STGNN architectures for this task. Table 2 provides a comparative overview of these models, emphasizing their key innovations, advantages, and limitations. In addition, LSTM and GRU models were implemented as temporal-only baselines, processing only temporal patterns without capturing spatial interactions between stations [46].

Table 2: Comparison of the four investigated STGNN models for forecasting.

Model	**Key innovation**	**Pros**	**Cons**
ASTGCN	Combines graph convolution, temporal convolution, and attention mechanisms to capture complex spatial and temporal dependencies [47].	High accuracy in modeling dynamic spatio-temporal data; effectively captures both spatial and temporal correlations.	High computational cost due to attention; limited scalability on very large networks.
STCONV	Streamlined integration of graph and temporal convolutions for efficient forecasting [48].	Computationally efficient; suitable for smaller datasets; low memory usage.	Limited ability to capture complex patterns in dynamic data.
MTGNN	Introduces multi-scale attention to capture dependencies at various time scales, ideal for multivariate data [49].	Strong performance on multivariate time-series forecasting; good for capturing long-term dependencies.	Complex architecture increases training time and resource demands; moderate scalability.
A3TGCN	Utilizes a 3D attention mechanism to model long-range spatial and temporal dependencies [50].	Effective for large-scale data; captures long-range dependencies in both spatial and temporal domains.	High computational cost; increased complexity can lead to higher resource requirements.

Fig. 3 presents the overall model framework. The spatio-temporal models take two types of input: temporal information, represented by a sliding window of time-series data, and spatial information, represented by a weighted adjacency matrix encoding inter-station relationships. Node feature vectors and the adjacency matrix are then fed into one of the selected STGNN models, where the data is processed through three layers of graph convolution to improve generalization, accuracy, and stability in graph-based learning.

In contrast, the baseline models employ LSTM and GRU architectures. In these time-series approaches, node feature vectors are processed independently through LSTM or GRU layers, capturing only temporal dependencies while ignoring spatial interactions. The outputs of the temporal layers are then passed through a fully connected layer, which aligns the predictions with the structure of the charging stations.

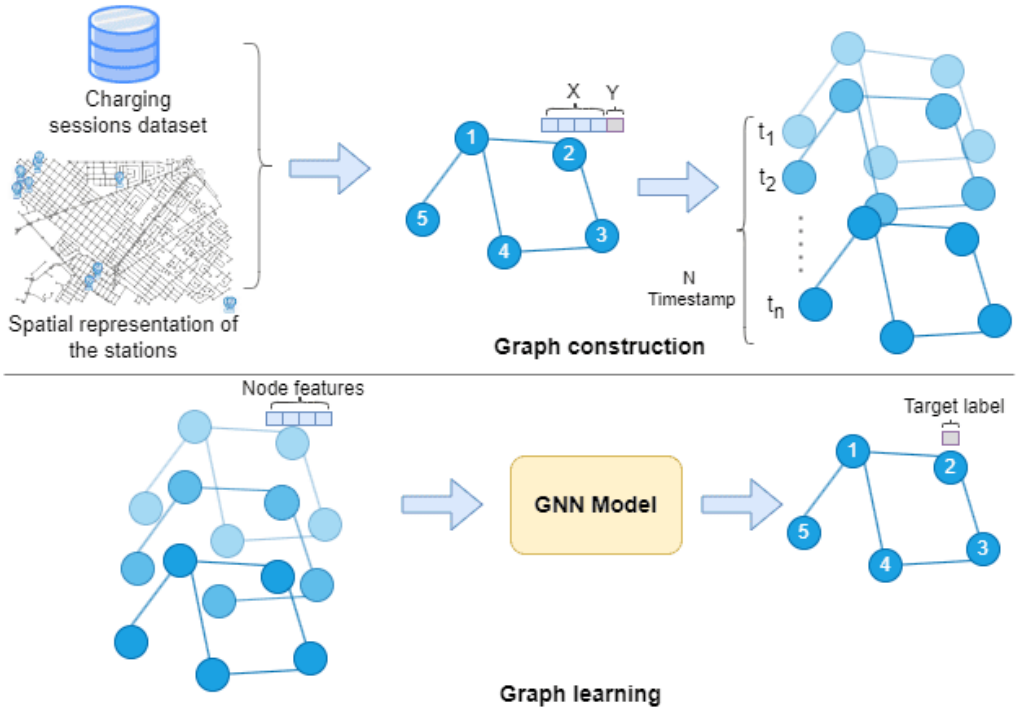


Figure 3: Proposed framework for forecasting energy demand at CSs.

4. Datasets

This section presents the two real-world datasets used in this study and outlines the data cleaning, smoothing-based noise reduction, and graph representation procedures applied for model implementation.

4.1. Datasets Overview

In this study, we utilize two datasets: the Yonne Administrative Department EVCS Dataset from France and the Palo Alto City EVCS Dataset from the United States. By incorporating datasets from two distinct geographic and socioeconomic contexts, we can better understand the model's adaptability to different station usage patterns and charging behaviors. This approach allows us to validate the STGNN's performance across varied real-world scenarios, comparing its predictive accuracy to diverse regions.

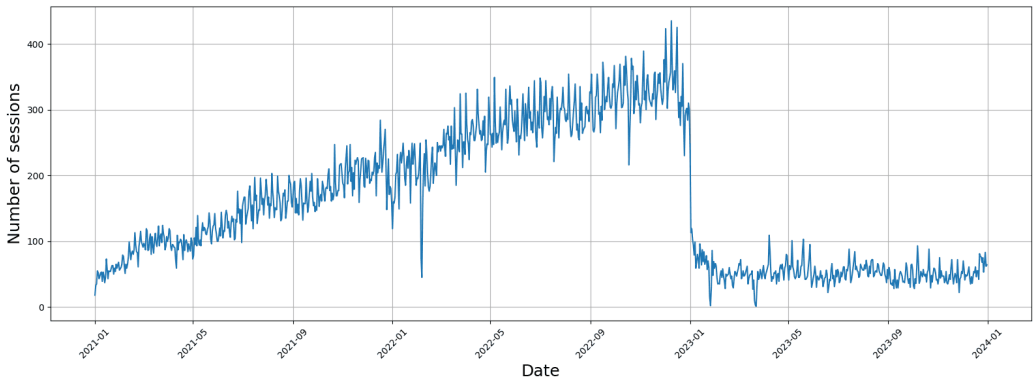
- **SDEY dataset:** The Yonne Administrative Department EVCS Dataset, provided by SDEY in France, contains charging session data spanning January 2021 to December 2023. This dataset covers over 133 charging points categorized into three types based on their charging capacity: normal ($\leq$25 kW), rapid (>25 kW), and ultra-rapid ($\geq$100 kW) types.
- **Palo Alto dataset:** The Palo Alto City EVCS Dataset provides extensive data on charging sessions collected from 46 CSs across Palo Alto, California, covering a period of over ten years (2011–2021), documenting approximately 259,376 charging sessions registered between 2011 and 2020. Notably, only 5 stations were established in 2011. Subsequently, 6 additional stations were installed in 2014, with a consistent expansion of 5 new stations per year from 2015 to 2018.

Both datasets provide detailed information on individual charging sessions, including start and end times, energy consumption, and charging speed. They constitute a valuable resource for analyzing EV charging behavior over extended periods and across different geographic contexts.

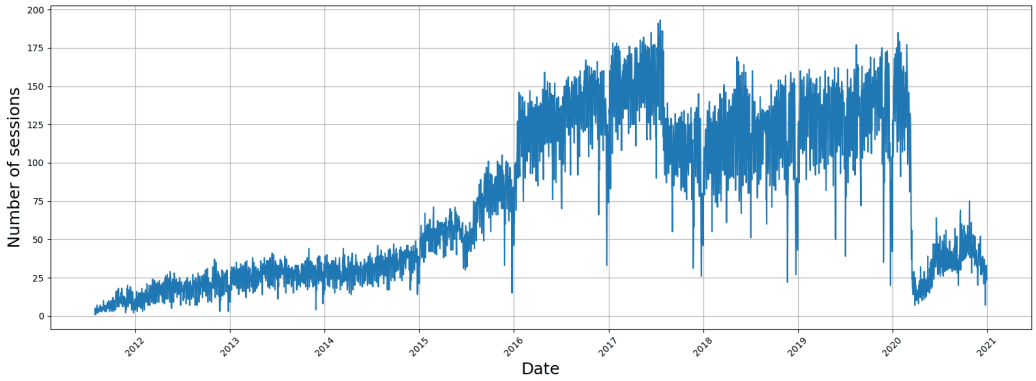
4.2. Data Analysis

The objective is to accurately forecast energy consumption at CSs. Figure 4a shows station usage for the SDEY dataset from January 2021 to December 2023. The results reveal steady growth in charging sessions with

a peak in 2021, reflecting gradual adoption and possible infrastructure expansion. A sharp decline in energy consumption is observed after the implementation of a pricing policy in January 2023, whereas the stations had previously been free to use. Similarly, Figure 4b shows usage patterns in Palo Alto across all stations from 2012 to early 2021. The data indicate a steady increase in charging sessions, with accelerated growth beginning around 2016, likely driven by rising EV adoption and the expansion of charging infrastructure. This trend peaks between 2017 and 2018, followed by fluctuations in daily session counts. A sharp decline in early 2020 coincides with the onset of the COVID-19 pandemic, which disrupted mobility and reduced station utilization.



(a) SDEY Dataset



(b) Palo Alto Dataset

Figure 4: Count of charging sessions over time for the SDEY and Palo Alto datasets.

Figures 4a and 4b highlight a key challenge in analyzing energy consumption data: the presence of substantial noise. Frequent fluctuations, irregular amplitudes, and sudden spikes obscure clear trends and complicate interpretation. To improve predictive accuracy, smoothing techniques (see Section 4.5) are applied to reduce noise, emphasize underlying patterns, and mitigate the impact of transient anomalies.

4.3. Data Preprocessing

The preprocessing of the datasets includes several steps to ensure data quality and consistency for reliable model training and analysis. Missing values were handled by removing records lacking essential information, such as unknown CS identifiers or absent energy consumption values, thereby preserving dataset integrity. Outliers were manually corrected, particularly cases of incorrect CS names or locations caused by data entry errors, ensuring that all records accurately represented real stations.

To standardize the datasets, energy consumption was aggregated on a daily basis, emphasizing overarching trends rather than fine-grained fluctuations. Temporal consistency was ensured by excluding data before each station's activation date, so only operational stations were included. Data from the COVID-19 period were also removed due to disruptions in usage, as many stations were inactive or exhibited atypical patterns. After preprocessing, the Palo Alto dataset spans January 2018 to March 2020, capturing regular pre-pandemic charging behavior, while the SDEY dataset covers January 2021 to December 2022, reflecting a high volume of sessions as noted in the previous subsection.

4.4. Features Engineering

Numerous studies have highlighted the influence of POIs on the utilization of EVCSs [51]. POIs play an important role in the strategic deployment of charging infrastructure. For example, shopping malls provide convenient charging opportunities for drivers while shopping, and highways offer essential charging locations for long-distance travelers. To enrich the datasets in this study, POI data were extracted from OpenStreetMap (OSM) and classified into relevant categories, as shown in Table 3. This procedure was applied consistently across both datasets, enabling the counting of POIs in each category within the vicinity of each CS and offering insights into the urban and contextual factors that may influence charging infrastructure usage.

Table 3: POI extracted from OSM.

POI Category	**Description**	**Count Palo Alto**	**Count SDEY**
Parking	Refers to parking lots or garages where vehicles can be parked and charge.	93	203
Residential	Includes residential areas like neighborhoods or apartment complexes, where home-based or nearby charging is common.	21	63
Office	Refers to offices and business complexes where EV owners may charge their vehicles while at work.	122	36
Primary	Major roads connecting important destinations, essential for long-distance EV travelers.	25	82
Secondary	Smaller roads connecting local regions, providing access for regional EV users.	73	151
Tertiary	Local roads connecting smaller towns, serving EV owners in suburban or rural areas.	179	169

Additional features were engineered to capture finer patterns in CS usage. Temporal variables included the day of the week, month, weekend, and holiday indicator, reflecting periodic consumption behaviors. Dynamic features, such as the number of sessions on the previous day, charging duration, and the trend component of energy demand, were also incorporated to enrich the datasets.

Furthermore, lagged features were incorporated to capture autocorrelation, where past consumption influences future demand. Specifically, Lagged segments include energy_lag_1, energy_lag_3, and energy_lag_7, representing usage from the previous day, three days earlier, and seven days earlier, respectively. In addition, rolling statistics were computed to smooth fluctuations and capture short-term trends. For instance, rolling_mean_3 and rolling_std_3 reflect the 3-day rolling average and standard deviation, while rolling_mean_7 and rolling_std_7 capture these statistics over a 7-

day window. Such features help represent local trends and volatility, improving model performance in forecasting future energy usage. Table 4 summarizes the engineered features used as model inputs.

Table 4: Summary of input features for the graph models.

Feature Category	Description
Station Information	Geographic coordinates (Latitude, Longitude), unique station identifier (ID_station)
Energy Data	Lagged energy values (1, 3, and 7 days), rolling mean and standard deviation (3 and 7 days), energy trend
Session and Duration Data	Session count and duration sum lagged by 1, 3, 7, 14, and 30 days
Location and Contextual Data	POIs near stations, station activity status (is_inactive), Population data
Temporal Data	Day of the week, month of the year, weekend and holiday indicators
Additional Features	Peak usage indicator (is_peak) and general trend indicator (trend)

4.5. Data Smoothing

Noisy data can significantly degrade the prediction performance of ML models. Noise can obscure underlying trends and reduce generalizability, leading to lower accuracy and reliability. Techniques such as Holt–Winters exponential smoothing and Exponential Moving Average (EMA) [52] are widely used to mitigate noise by stabilizing fluctuations and allowing models to focus on essential patterns. In this study, noise reduction is performed using EMA, a widely used smoothing method for time-series data. EMA assigns greater weight to recent observations while gradually reducing the influence of older values, making it responsive to recent changes. This approach smooths short-term fluctuations caused by random noise while preserving the underlying trend. The EMA at time t is defined as:

$$EWMA_t = \alpha \times X_t + (1 - \alpha) \times EMA_{t-1} \tag{9}$$

where X_t is the data point at time t, and $\alpha = \frac{2}{n+1}$ is the smoothing factor. A higher α makes EMA more sensitive to recent variations, whereas a lower α

produces smoother curves. Choosing an appropriate α is crucial to balance responsiveness and stability.

For illustration, Figure 5 shows the actual energy values alongside the smoothed series generated using $\alpha = 0.8$. The EMA yields a steadier curve by reducing spikes and dips while preserving overall trends. This enhances the visibility of long-term patterns, reduces short-term volatility, and provides a clearer representation of the data's core behavior.

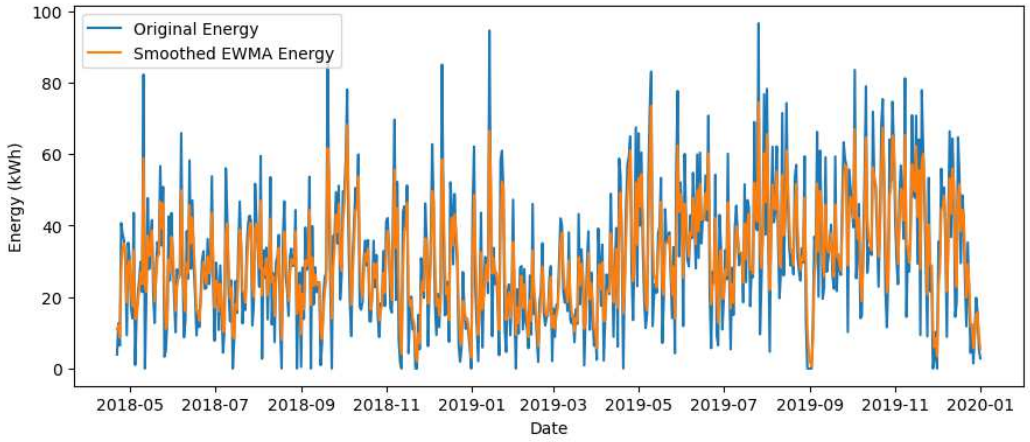


Figure 5: Comparison of original and smoothed sample data values from the SDEY dataset.

4.6. Graph Representation of Data

For the graph-based model, the data was represented as a graph with nodes (stations) and edges capturing spatial and temporal dependencies. The road network was extracted from OSM to create directed edges representing the shortest paths between stations, with distance used as the edge weight. A simplification step was applied to reduce redundant edges and eliminate duplicates, resulting in a streamlined graph that connects all stations and effectively encodes the required spatial and temporal information, as shown in Fig. 6.

5. Experiments and Results

This section presents the evaluation metrics, describes the experimental setup, and analyzes the results to assess the proposed framework against baseline models and demonstrate its effectiveness in forecasting energy demand within EVCS networks.

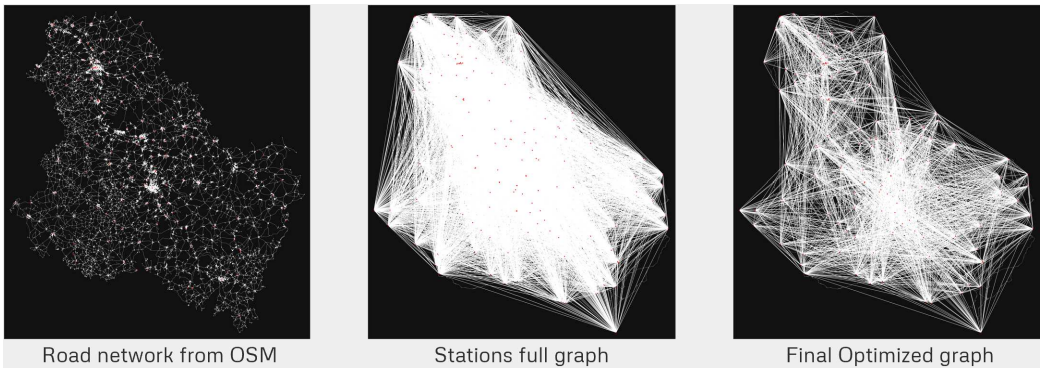


Figure 6: Preprocessing workflow for edge simplification on SDEY dataset.

5.1. Evaluation Metrics

To evaluate the effectiveness of the proposed approaches, several performance metrics were employed to assess the STGNN models, summarized as follows.

- Mean Absolute Error (MAE): calculates the average absolute difference between actual and predicted values. Lower MAE indicates better model performance; it is calculated by:

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i| \tag{10}$$

  Where y_i is the actual value, $\hat{y}_i$ is the predicted value, and n is the number of data points.

- Mean Squared Error (MSE): measures the average of the squared differences between actual and predicted values. A lower MSE is preferred, implying the model's predictions are closer to the actual values, especially minimizing larger errors. The formula for MSE is:

$$\text{MSE} = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2 \tag{11}$$

- Root Mean Squared Error (RMSE): derived from MSE, it measures the magnitude of prediction errors. Lower RMSE values indicate higher

accuracy and closer alignment with actual values. It's calculated by:

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{12}$$

- Coefficient of Determination (R^2): evaluates model fit, ranging from 0 to 1, with higher values indicating a better fit. An R^2 close to 1 means the model explains most of the variance in the data. The formula for R^2 is:

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2} \tag{13}$$

  where $\bar{y}$ is the mean of the actual values.

5.2. Forecasting Models

In this study, the data was aggregated on a daily basis, and each input sequence consisted of 7 consecutive days, used to predict the energy demand for the following day. This approach ensures that the model learns temporal patterns in the data. We proposed four GNN architectures: MTGNN, A3TGCN, STCONV and the ASTGCN. For each station node $v \in V$, these models use a time-series input of historical energy data over 7 days to predict the energy demand for the next day. The input of node v is represented as a sequence of feature vector $\mathbf{X}_v = \{X_v^{(t-7)}, X_v^{(t-6)}, \ldots, X_v^{(t-1)}\}$.

Figure 7 illustrates the workflow of the proposed energy demand prediction system for EVCSs. The process begins with data preprocessing and engineering, where raw data is collected from CSs and temporal and spatial information are extracted, as explained in Section 4. After the graph construction phase, the data is split into 60% for training, 20% for validation, and 20% for testing. During the training and validation phase, one spatio-temporal model is selected. This is followed by the application of multiple convolutional layers (Conv2D, BatchNorm, and ReLU) to learn data representations, with training conducted using the hyperparameters provided in Table 5.

Once the model is trained and validated using the validation data, it is evaluated in the testing phase using the 20% test data to generate daily energy demand predictions. Finally, performance metrics such are computed to assess the accuracy and effectiveness of the trained model.

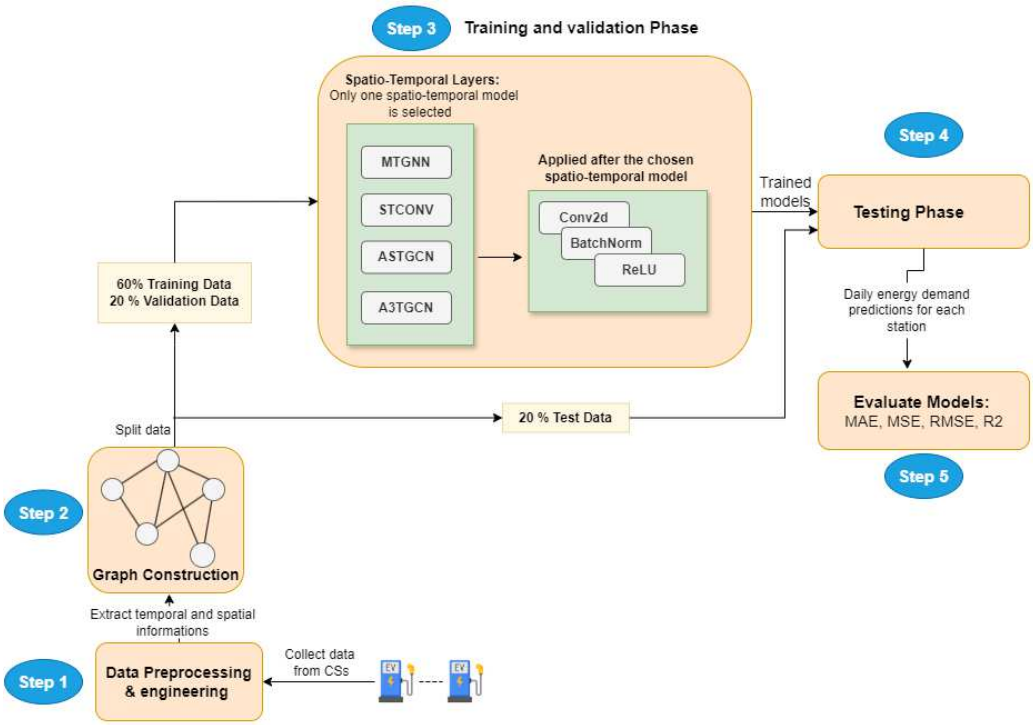


Figure 7: Overview of the architecture of the graph-based proposed models.

Table 5: Hyperparameter configuration for all the models.

Hyperparameter	MTGNN	ASTGCN	STCONV	A3TGCN	LSTM	GRU
Epochs	150	120	100	200	80	90
Weight Decay	1e−3	5e−4	1e−4	5e−3	1e−2	1e−3
Common Settings:						
Learning Rate: 0.01	Batch Size: 64	Scheduler: StepLR	Step Size: 30	Gamma: 0.1	Optimizer: Adam	

In addition to the GNN-based models, we developed two baseline models using LSTM and GRU networks. These models, widely recognized for time-series forecasting, utilize the same input format X_v to predict Y_v^t but treat each station independently, without incorporating graph-based spatial relationships. This setup provides a comparison to assess the impact of the GNN architecture on predictive accuracy and demonstrate the advantages of leveraging spatial dependencies within the charging network.

Table 6 summarizes the performance metrics for graph-based and baseline models on the Palo Alto dataset. The results show that MTGNN delivers

the best overall performance, achieving the lowest test MSE (69.15), RMSE (8.32), and MAE (6.12), while maintaining the highest R^2 score (0.86) and an execution time of just 0.12 minutes. MTGNN's superior performance can be attributed to its ability to handle multivariate time series data more effectively by leveraging an adaptive graph learning mechanism. Unlike other models, MTGNN use of temporal convolutions further enhances its capacity to process multiple time-dependent variables efficiently, making it particularly suitable for large-scale energy forecasting tasks.

Table 6: Performance metrics for graph-based and baseline models on the Palo Alto dataset by phase.

		Graph-based Models				**Baseline Models**	
		MTGNN	**ASTGCN**	**STCONV**	**A3TGCN**	**LSTM**	**GRU**
Train	MSE	**46.41**	51.24	51.59	70.16	95.60	119.03
	RMSE	**6.81**	7.16	7.18	8.38	9.78	10.91
	MAE	**5.16**	5.40	5.38	6.30	7.41	8.52
	R^2	**0.87**	0.86	0.86	0.81	0.74	0.68
Validation	MSE	**63.51**	74.32	73.18	110.68	99.08	103.03
	RMSE	**7.97**	8.62	8.55	10.52	9.95	10.15
	MAE	**5.97**	6.47	6.37	7.79	8.01	8.15
	R^2	**0.85**	0.82	0.83	0.74	0.78	0.77
Test	MSE	**69.15**	81.30	90.89	137.67	113.34	127.90
	RMSE	**8.32**	9.02	9.53	11.73	10.65	11.31
	MAE	**6.12**	6.73	7.05	8.79	8.43	8.92
	R^2	**0.86**	0.84	0.82	0.73	0.79	0.76
Execution Time (min)		**0.12**	1.56	58.79	0.09	23.13	20.86

ASTGCN also performs well, achieving the second-best accuracy with a test MSE of 81.30 and an R^2 score of 0.84. ASTGCN benefits from its spatio-temporal attention mechanisms, which enable it to effectively capture spatial dependencies and temporal patterns. While it trails MTGNN in accuracy, it balances accuracy and computational efficiency with an execution time of 1.56 minutes, significantly faster than STCONV but slower than A3TGCN. These results suggest that ASTGCN is a robust choice for scenarios prioritizing accuracy with moderate computational resources.

STCONV, while accurate with a test MSE of 90.89 and an R^2 score of 0.82, suffers from an exceptionally long execution time of 58.79 minutes.

This indicates that STCONV's reliance on stacked spatial and temporal convolutional layers introduces significant computational overhead. Although it performs better than A3TGCN in terms of accuracy, its inefficiency in execution time limits its practicality for large-scale, real-time scenarios.

A3TGCN offers moderate accuracy improvements with a test MSE of 137.67 and an R^2 score of 0.73. Its use of attention mechanisms enables it to capture complex dependencies, but its generalization ability is hindered during periods of high variability in energy demand, leading to reduced performance compared to MTGNN and ASTGCN. However, A3TGCN is highly efficient in terms of execution time, completing predictions in only 0.09 minutes, which is even faster than MTGNN. This makes A3TGCN a viable option for applications requiring rapid predictions, provided the variability in the data is low.

Baseline models such as LSTM (MSE = 113.34, R^2 = 0.79) and GRU (MSE = 127.90, R^2 = 0.76) demonstrate limited accuracy due to their inability to model spatial dependencies. These models focus solely on temporal patterns, which is insufficient for capturing the complex relationships inherent in multivariate spatiotemporal data like EVCS energy demand. Additionally, their execution times (23.13 and 20.86 minutes, respectively) are significantly longer than MTGNN and ASTGCN, further underscoring their inefficiency for this task.

Table 7 presents the performance metrics for forecasting one-day-ahead energy consumption using the SDEY dataset across four STGNN models: MTGNN, ASTGCN, STCONV, and A3TGCN, as described in Section 3.3. Additionally, the table includes results from two baseline models, LSTM and GRU.

For the SDEY dataset, graph-based models consistently exhibit superior performance compared to the baseline models. Among these, the ASTGCN model achieves the highest accuracy, reflected in a test RMSE of 15.16 and an R^2 score of 0.86. Its ability to effectively model spatial and temporal dependencies using attention mechanisms contributes to its superior predictive power. Closely following, the MTGNN model demonstrates strong performance with a test RMSE of 15.54, a minor difference compared to ASTGCN. MTGNN also achieves an R^2 score of 0.86, indicating its effectiveness in capturing demand variability across different phases of the dataset. The small gap between ASTGCN and MTGNN underscores the reliability of both models for energy demand forecasting in complex datasets.

Table 7: Performance metrics for graph-based and baseline models for the SDEY dataset by phase.

		Graph-based Models				Baseline Models	
		MTGNN	**ASTGCN**	**STCONV**	**A3TGCN**	**LSTM**	**GRU**
Train	MSE	135.23	**114.68**	139.03	232.66	336.77	361.53
	RMSE	11.63	**10.71**	11.80	15.25	18.35	19.01
	MAE	7.07	**6.85**	7.62	8.72	10.39	10.76
	R^2	0.86	**0.89**	0.86	0.77	0.70	0.68
Validation	MSE	231.86	**211.97**	298.62	543.90	313.84	379.55
	RMSE	15.23	**14.56**	17.28	23.32	17.72	19.48
	MAE	**9.35**	9.39	11.11	14.09	11.51	12.40
	R^2	0.85	**0.86**	0.80	0.65	0.81	0.77
Test	MSE	241.37	**229.79**	492.63	860.65	332.67	414.61
	RMSE	15.54	**15.16**	22.20	29.34	18.24	20.36
	MAE	10.27	**10.22**	14.74	19.73	12.39	13.49
	R^2	**0.86**	**0.86**	0.71	0.51	0.82	0.77
Execution Time (min)		**1.39**	2.14	47.43	0.67	73.03	109.39

STCONV shows good predictive power, with a test MSE of 492.63 and RMSE of 22.20. However, its execution time of 47.43 minutes is significantly longer than those of ASTGCN (2.14 minutes) and MTGNN (1.39 minutes). This computational overhead limits its feasibility for time-sensitive forecasting applications, despite its reasonable accuracy. In contrast, the A3TGCN model exhibits higher error rates, with a test MSE of 860.65 and RMSE of 29.34, and an R^2 score of 0.51, reflecting difficulties in generalizing across demand patterns with high variability. This suggests that A3TGCN struggles to adapt to datasets with complex temporal and spatial characteristics, making it less robust in this scenario.

The baseline models, LSTM and GRU, exhibit limited performance compared to graph-based models. LSTM achieves a test MSE of 332.67 and RMSE of 18.24, while GRU reports a test MSE of 414.61 and RMSE of 20.36. These results highlight the shortcomings of baseline models in capturing spatial dependencies, which are critical in the graph of CSs. Furthermore, their longer execution times (73.03 minutes for LSTM and 109.39 minutes for GRU) emphasize their inefficiency relative to STGNN models, further limiting their practicality for real-time or large-scale forecasting tasks.

The differences in performance across training, validation, and test sets

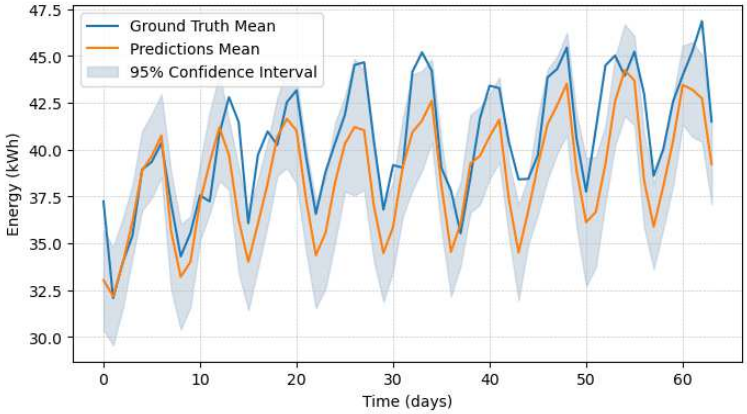
are influenced by the time-based data splits, where data is split chronologically. This approach ensures realistic forecasting but may result in slight variations in station usage across phases, as seen in the SDEY dataset.

The execution time of 0.12 minutes for the Palo Alto dataset and 1.39 minutes for the SDEY dataset highlights MTGNN's computational efficiency, making it well-suited for applications that demand timely forecasts. The higher execution time for the SDEY dataset is attributed to its greater complexity, with 130 station nodes compared to 33 in the Palo Alto dataset. Despite this difference, MTGNN maintains competitive performance, demonstrating its scalability to larger and more complex datasets.

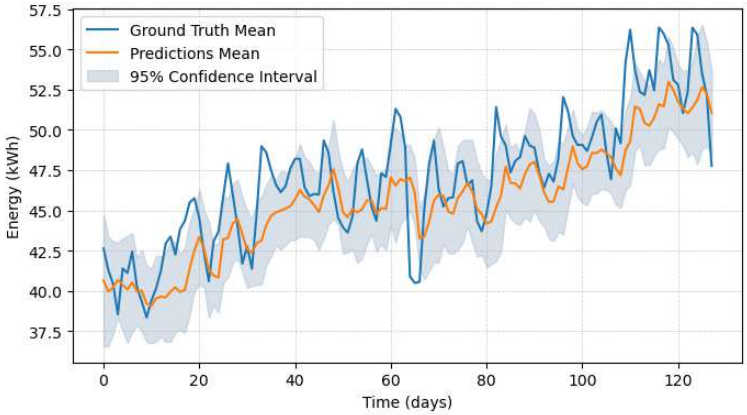
In summary, the results confirm that graph-based models, particularly MTGNN and ASTGCN, consistently outperform baseline models (LSTM and GRU) across diverse datasets from different regions, such as France and the USA. These findings highlight the value of spatio-temporal modeling for energy demand forecasting in EVCS networks and demonstrate MTGNN's robustness and adaptability across varied datasets and complexities.

Fig. 8 illustrates the mean energy consumption predictions for all CSs in the Palo Alto dataset (Fig. 8a) and SDEY dataset (Fig. 8b), with a 95% confidence interval shaded in grey to highlight the MTGNN's predictive uncertainty. We calculated the average forecast across all CSs, providing a comprehensive view of the MTGNN model's performance. The model's predictions (orange line) align strongly with the actual observed values (blue line), particularly in capturing both datasets' general trends and seasonal fluctuations. The Palo Alto dataset displays more consistent seasonal peaks and smoother demand transitions, resulting in narrower confidence intervals and slightly closer alignment between predictions and actual values. However, in the SDEY dataset, the model encounters more significant fluctuations and wider confidence intervals, particularly around peak values. This dataset exhibits higher variability and more irregular demand patterns, which increases forecasting complexity and uncertainty. Overall, the model's performance across both datasets underscores its robustness and adaptability in forecasting daily energy consumption across geographically and behaviorally diverse EVCS networks, with acceptable uncertainty quantification in stable and highly variable demand scenarios.

Fig. 9 presents the residuals (or the differences between the actual and predicted energy consumption values) for the two datasets using the MTGNN model on the test set. The residuals are centered around zero in both figures,



(a) Palo Alto Dataset

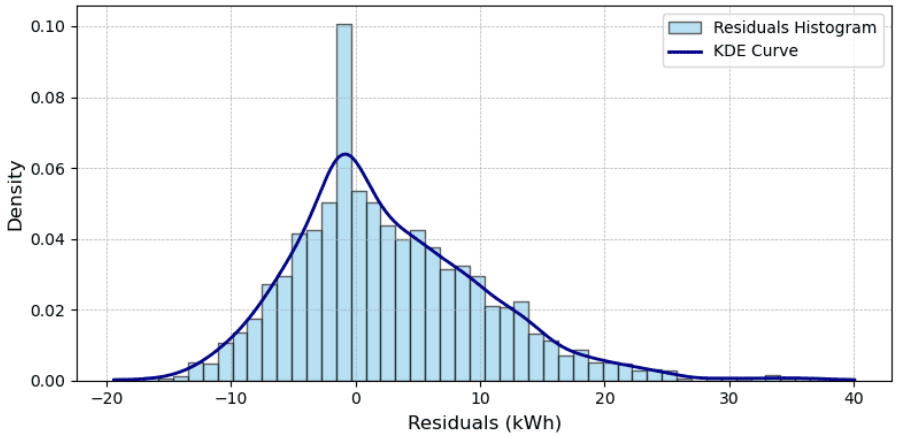


(b) SDEY Dataset

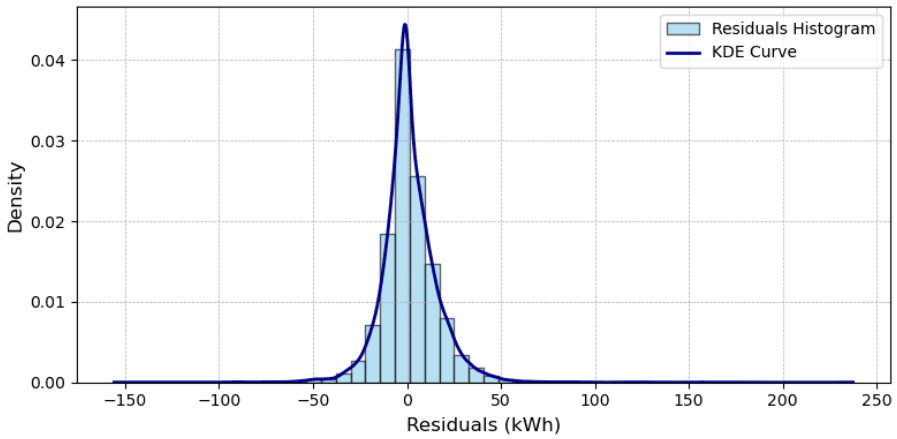
Figure 8: The mean predictions of all CSs for one day ahead on the test set using the MTGNN model.

indicating that the model generally produces unbiased predictions. However, notable differences in the spread and skew of the residuals highlight distinct characteristics of each dataset. The SDEY dataset (see Fig. 9b) has a few extreme residuals ranging from approximately -150 kWh to 250 kWh. This spread suggests that the MTGNN model experiences larger deviations in its predictions, likely due to the higher variability in energy consumption at certain stations within this dataset. In contrast, the residuals for the Palo Alto dataset (see Fig. 9a) are mostly contained within a narrower range

of -20 kWh to 40 kWh, suggesting that the model achieves more consistent accuracy with fewer extreme deviations. This tighter distribution indicates that energy consumption patterns in Palo Alto are more stable, allowing the MTGNN model to deliver more reliable predictions for this dataset.



(a) Palo Alto Dataset



(b) SDEY Dataset

Figure 9: Distribution of residuals across all stations in the test set using the MTGNN model.

6. Conclusion

Charging infrastructure for EVs is crucial in driving sustainable transportation forward and reducing carbon emissions. Enhancing EVCS infrastructure is essential to meet the rising demand, encourage EV adoption, and ensure efficient station utilization. This research presents a framework for energy demand forecasting. In this work, we developed and implemented four STGNN models to accurately forecast daily energy demand by capturing spatial and temporal patterns in station usage. We compared these models to baseline approaches such as LSTM and GRU using two diverse datasets from the US and France. Through extensive evaluation, we identified the MTGNN model as the top performer, achieving an R^2 of 86% on both datasets and demonstrating the fastest execution time, making it well-suited for accurate and efficient energy demand forecasting. Future work will focus on enhancing long-term forecasting accuracy by integrating additional data sources, such as real-time traffic information, and refining the optimization process using reinforcement learning. From a methodological perspective, two promising directions stand out. First, transfer learning approaches can adapt models trained in data-rich regions (e.g., urban areas with dense charging infrastructure) to data-scarce contexts, reducing the need for extensive retraining while maintaining accuracy. This is particularly relevant for emerging markets and semi-rural areas where charging data is limited but infrastructure deployment is accelerating. Second, meta-learning techniques offer the potential for rapid model adaptation to new CSs or evolving demand conditions. By learning "how to learn" across multiple tasks, meta-learning enables the model to generalize quickly with minimal additional data, making it highly suitable for dynamic and heterogeneous EVCS environments. These directions strengthen the scalability and real-world applicability of STGNN-based forecasting.

Acknowledgment

This paper is supported by the OPEVA project that has received funding within the Chips Joint Undertaking (Chips JU) from the European Union's Horizon Europe Programme and the National Authorities (France, Czechia, Italy, Portugal, Turkey, Switzerland), under grant agreement 101097267. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or Chips JU. This work was also supported by the SDEY (Syndicat Départemental d'Energies de l'Yonne).

References

[1] K. Shahzad and I. I. Cheema, "Low-carbon technologies in automotive industry and decarbonizing transport," *Journal of Power Sources*, vol. 591, p. 233888, 2024.

[2] A. Tayri and X. Ma, "Grid impacts of electric vehicle charging: A review of challenges and mitigation strategies," *Energies*, vol. 18, no. 14, p. 3807, 2025.

[3] International Energy Agency, "Outlook for energy demand," https://www.iea.org/reports/global-ev-outlook-2025/outlook-for-energy-demand, 2025, accessed: 2025-08-18. [Online]. Available: https://www.iea.org/reports/global-ev-outlook-2025/outlook-for-energy-demand

[4] A. Orzechowski, L. Lugosch, H. Shu, R. Yang, W. Li, and B. H. Meyer, "A data-driven framework for medium-term electric vehicle charging demand forecasting," *Energy and AI*, vol. 14, p. 100267, 2023.

[5] T. Zhang, Q. Peng, and S. Zeng, "Predicting ev charging demand in renewable-energy-powered grids using explainable machine learning," *Sustainability*, vol. 17, no. 9, p. 4158, 2025.

[6] F. Hao, "Impact of electric vehicle charging demand on clean energy regional power grid control," *Energy Informatics*, vol. 8, no. 1, p. 83, 2025.

[7] A. Mystakidis, N. Tsalikidis, P. Koukaras, G. Skaltsis, D. Ioannidis, C. Tjortjis, and D. Tzovaras, "Ev charging forecasting exploiting traffic, weather and user information," *International Journal of Machine Learning and Cybernetics*, pp. 1–27, 2025.

[8] T. Aldhanhani, A. Abraham, W. Hamidouche, and M. Shaaban, "Future trends in smart green iov: Vehicle-to-everything in the era of electric vehicles," *IEEE Open Journal of Vehicular Technology*, vol. 5, pp. 278–297, 2024.

[9] M. Rashid, T. Elfouly, and N. Chen, "A comprehensive survey of electric vehicle charging demand forecasting techniques," *IEEE Open Journal of Vehicular Technology*, 2024.

[10] H. Liang, I. Sharma, W. Zhuang, and K. Bhattacharya, "Plug-in electric vehicle charging demand estimation based on queueing network analysis," in *2014 IEEE PES General Meeting| Conference & Exposition*. IEEE, 2014, pp. 1–5.

[11] R. N. Akshayaa, S. Sajidha, and R. Radha, "Session and energy forecasting for electric vehicle charging station using custom weighted ensemble machine learning," *Results in Engineering*, p. 105104, 2025.

[12] L. Douaidi, S.-M. Senouci, I. El Korbi, F. Harrou, and A. Yazıcı, "Federated deep learning for enhanced prediction of electric vehicle charging station availability," *Cluster Computing*, vol. 28, no. 6, p. 364, 2025.

[13] M. W. Ali, A. b. Mustafa, M. A. M. Shuvo, and B. Sick, "Location based probabilistic load forecasting of ev charging sites: Deep transfer learning with multi-quantile temporal convolutional network," *arXiv preprint arXiv:2409.11862*, 2024.

[14] Y. Zhang, C. Liu, X. Rao, X. Zhang, and Y. Zhou, "Spatial-temporal load forecasting of electric vehicle charging stations based on graph neural network," *Journal of Intelligent & Fuzzy Systems*, vol. 46, no. 1, pp. 821–836, 2024.

[15] A. Palaniappan, P. Bhukya, S. K. Chitti, and J. Gao, "Data-driven analysis of ev energy prediction and planning of ev charging infrastructure," in *2023 IEEE Ninth International Conference on Big Data Computing Service and Applications (BigDataService)*. IEEE, 2023, pp. 17–24.

[16] K. H. Baesmat, F. Shokoohi, and Z. Farrokhi, "Sp-rf-arima: A sparse random forest and arima hybrid model for electric load forecasting," *Global Energy Interconnection*, vol. 8, no. 3, pp. 486–496, 2025.

[17] S. Wu, Y. Xiao, S. Fu, J. Choi, and C. Zheng, "A hybrid deep learning model for load forecasting of electric vehicle charging stations using time series decomposition," *Journal of Power Sources*, vol. 655, p. 237882, 2025.

[18] 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," in *2024 6th Asia Energy and Electrical Engineering Symposium (AEEES)*. IEEE, 2024, pp. 1180–1185.

[19] K. Zheng, H. Xu, Z. Long, Y. Wang, and Q. Chen, "Coherent hierarchical probabilistic forecasting of electric vehicle charging demand," *IEEE Transactions on Industry Applications*, 2023.

[20] S. Wang, A. Chen, P. Wang, and C. Zhuge, "Predicting electric vehicle charging demand using a heterogeneous spatio-temporal graph convolutional network," *Transportation Research Part C: Emerging Technologies*, vol. 153, p. 104205, 2023.

[21] F. B. Hüttel, I. Peled, F. Rodrigues, and F. C. Pereira, "Deep spatio-temporal forecasting of electrical vehicle charging demand," *arXiv preprint arXiv:2106.10940*, 2021.

[22] A.-R. Houalef, F. Delavernhe, S.-M. Senouci, and E.-H. Aglzim, "Utilizing data-driven techniques to improve predictive modeling of connected electric vehicle energy consumption," in *2024 IEEE 100th Vehicular Technology Conference (VTC2024-Fall)*. IEEE, 2024, pp. 1–6.

[23] F. Rassaei, W.-S. Soh, and K.-C. Chua, "A statistical modelling and analysis of residential electric vehicles' charging demand in smart grids," in *2015 IEEE power & energy society innovative smart grid technologies conference (ISGT)*. IEEE, 2015, pp. 1–5.

[24] P. Aduama, Z. Zhang, and A. S. Al-Sumaiti, "Multi-feature data fusion-based load forecasting of electric vehicle charging stations using a deep learning model," *Energies*, vol. 16, no. 3, p. 1309, 2023.

[25] A. Ostermann and T. Haug, "Probabilistic forecast of electric vehicle charging demand: analysis of different aggregation levels and energy procurement," *Energy Informatics*, vol. 7, no. 1, p. 13, 2024.

[26] F. Ren, C. Tian, G. Zhang, C. Li, and Y. Zhai, "A hybrid method for power demand prediction of electric vehicles based on sarima and deep learning with integration of periodic features," *Energy*, vol. 250, p. 123738, 2022.

[27] E. Dokur, N. Erdogan, and S. Kucuksari, "Ev fleet charging load forecasting based on multiple decomposition with ceemdan and swarm decomposition," *IEEE access*, vol. 10, pp. 62 330–62 340, 2022.

[28] K. Lee and W. Rhee, "Ddp-gcn: Multi-graph convolutional network for spatiotemporal traffic forecasting," *Transportation Research Part C: Emerging Technologies*, vol. 134, p. 103466, 2022.

[29] P. Banda, M. A. Bhuiyan, K. N. Hasan, and K. Zhang, "Assessment of hybrid transfer learning method for forecasting ev profile and system voltage using limited ev charging data," *Sustainable Energy, Grids and Networks*, vol. 36, p. 101191, 2023.

[30] M. Cavus, H. Ayan, D. Dissanayake, A. Sharma, S. Deb, and M. Bell, "Forecasting electric vehicle charging demand in smart cities using hybrid deep learning of regional spatial behaviours," *Energies*, vol. 18, no. 13, p. 3425, 2025.

[31] C. Peng, C. Xu, and L. Jiao, "Variational spatiotemporal factorized graph neural networks for integrated electric vehicle usage prediction and pattern recognition with missing data on regional networks," *Energy*, p. 137821, 2025.

[32] N. Huang, Q. He, J. Qi, Q. Hu, R. Wang, G. Cai, and D. Yang, "Multinodes interval electric vehicle day-ahead charging load forecasting based on joint adversarial generation," *International Journal of Electrical Power & Energy Systems*, vol. 143, p. 108404, 2022.

[33] H. Jahangir, S. S. Gougheri, B. Vatandoust, M. A. Golkar, M. A. Golkar, A. Ahmadian, and A. Hajizadeh, "A novel cross-case electric vehicle demand modeling based on 3d convolutional generative adversarial networks," *IEEE Transactions on Power Systems*, vol. 37, no. 2, pp. 1173–1183, 2021.

[34] H. J. Kim and M. K. Kim, "Spatial-temporal graph convolutional-based recurrent network for electric vehicle charging stations demand forecasting in energy market," *IEEE Transactions on Smart Grid*, 2024.

[35] S. R. Fahim, R. Atat, C. Kececi, A. Takiddin, M. Ismail, K. R. Davis, and E. Serpedin, "Forecasting ev charging demand: A graph convolutional neural network-based approach," in *2024 4th International Conference on Smart Grid and Renewable Energy (SGRE)*. IEEE, 2024, pp. 1–6.

[36] J. Shi, W. Zhang, Y. Bao, D. W. Gao, and Z. Wang, "Load forecasting of electric vehicle charging stations: Attention based spatiotemporal multigraph convolutional networks," *IEEE Transactions on Smart Grid*, 2023.

[37] J. Tang, L. Xia, and C. Huang, "Explainable spatio-temporal graph neural networks," in *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 2023, pp. 2432–2441.

[38] B. Yu, H. Yin, and Z. Zhu, "Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting," *arXiv preprint arXiv:1709.04875*, 2017.

[39] Z. Wu, S. Pan, G. Long, J. Jiang, X. Chang, and C. Zhang, "Connecting the dots: Multivariate time series forecasting with graph neural networks," in *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, 2020, pp. 753–763.

[40] Y. Li, R. Yu, C. Shahabi, and Y. Liu, "Diffusion convolutional recurrent neural network: Data-driven traffic forecasting," *arXiv preprint arXiv:1707.01926*, 2017.

[41] Z. Fang, Q. Long, G. Song, and K. Xie, "Spatial-temporal graph ode networks for traffic flow forecasting," in *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, 2021, pp. 364–373.

[42] I. E. Agency, "Global ev outlook 2023," 2023, accessed: 28/07/2024. [Online]. Available: https://www.iea.org/reports/global-ev-outlook-2023/trends-in-charging-infrastructure

[43] Z. A. Sahili and M. Awad, "Spatio-temporal graph neural networks: A survey," *arXiv preprint arXiv:2301.10569*, 2023.

[44] M. Newman, "Message passing methods on complex networks," *Proceedings of the Royal Society A*, vol. 479, no. 2270, p. 20220774, 2023.

[45] M. Jin, H. Y. Koh, Q. Wen, D. Zambon, C. Alippi, G. I. Webb, I. King, and S. Pan, "A survey on graph neural networks for time series: Forecasting, classification, imputation, and anomaly detection," *arXiv preprint arXiv:2307.03759*, 2023.

[46] F. M. Shiri, T. Perumal, N. Mustapha, and R. Mohamed, "A comprehensive overview and comparative analysis on deep learning models: Cnn, rnn, lstm, gru," *arXiv preprint arXiv:2305.17473*, 2023.

[47] S. Yan, Y. Xiong, and D. Lin, "Spatial temporal graph convolutional networks for skeleton-based action recognition," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 32, 2018.

[48] Z. A. Sahili and M. Awad, "Spatio-temporal graph neural networks: A survey," *arXiv preprint arXiv:2301.10569*, 2023.

[49] Z. Wu, S. Pan, G. Long, J. Jiang, X. Chang, and C. Zhang, "Connecting the dots: Multivariate time series forecasting with graph neural networks," in *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, 2020, pp. 753–763.

[50] J. Bai, J. Zhu, Y. Song, L. Zhao, Z. Hou, R. Du, and H. Li, "A3t-gcn: Attention temporal graph convolutional network for traffic forecasting," *ISPRS International Journal of Geo-Information*, vol. 10, no. 7, p. 485, 2021.

[51] V. Radermecker and L. Vanhaverbeke, "Estimation of public charging demand using cellphone data and points of interest-based segmentation," *World Electric Vehicle Journal*, vol. 14, no. 2, p. 35, 2023.

[52] F. Harrou, A. Zeroual, F. Kadri, and Y. Sun, "Enhancing road traffic flow prediction with improved deep learning using wavelet transforms," *Results in Engineering*, p. 102342, 2024.