International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

STGATE: An Explainable Spatio-Temporal Graph Attention Network for Multi-Task EV Charging Analytics

Sakshi Limaye, Shravani Nagawade, Jay Paspule
Ms. Divya Mary Biji
Assistant Professor, Department of Computer Science
Ajeenkya DY Patil University
Charholi bk, Pune, Maharashtra, India

Abstract— The rapid rise of electric vehicles (EVs) has created a growing need for smarter charging infrastructure. Most existing methods treat demand forecasting, user behavior, and grid monitoring as separate problems, even though they are closely connected in real-world usage. In this work, we propose STGATE (Spatio-Temporal Graph Attention Network with Transformer Encoder), a unified framework that addresses these challenges together. The model represents charging stations as nodes in a graph and uses a two-layer Graph Attention Network with dual attention heads to capture spatial relationships. At the same time, a two-layer Transformer Encoder learns temporal patterns from historical charging data. These spatial and temporal features are combined using a fusion layer and passed through a fully connected network with BatchNorm, ReLU, and Dropout to support multiple tasks, including demand forecasting, behavior clustering, grid stress prediction, and anomaly detection. To improve interpretability, the model includes SHAP-based feature analysis, node importance evaluation, and attention visualization. We evaluated the system on 1,320 charging sessions across 462 stations in five U.S. cities, achieving strong results, including an R^2 of 0.9494. The system is deployed using FastAPI with a React dashboard for real-time monitoring.

Keywords— *Electric Vehicles, Spatio-Temporal Graph Attention Network, Demand Forecasting, Anomaly Detection, Smart Grid, Deep Learning.*

I. INTRODUCTION

The rapid growth of electric vehicles (EVs) worldwide is transforming urban energy systems at an unprecedented pace. Global EV adoption crossed 40 million vehicles in 2023 and is projected to reach over 145 million by 2030 [1]. While this growth is encouraging for sustainability and reduction of carbon emissions, it also places increasing pressure on existing charging infrastructure and energy distribution networks. As EV usage expands in urban areas, charging networks are beginning to face issues such as peak demand spikes, uneven station utilization, and highly variable user behavior patterns across different regions and time periods [2].

Most traditional approaches to managing EV charging systems treat problems like energy demand forecasting, user behavior analysis, and grid monitoring as separate and independent tasks. In practice, however, these aspects are deeply interconnected, and separating them often leads to inefficient and fragmented solutions that fail to capture spatial and temporal dependencies between charging stations, ultimately reducing system efficiency and reliability [3].

To address this limitation, recent research has explored graph-based deep learning models, which are particularly effective in modeling spatial relationships across distributed charging stations within a networked environment [4]. When combined with temporal models, these approaches can also capture how charging patterns evolve dynamically over time under varying conditions [5]. Machine learning techniques have shown clear advantages over traditional statistical methods, particularly in short-term forecasting, adaptive learning, and anomaly detection tasks [6].

However, most existing solutions still focus on individual objectives rather than providing a unified and scalable framework. In addition, many deep learning models lack interpretability, making them less useful for real-world decision-making by infrastructure planners and policymakers who require transparency and trust in predictions [7][8].

In this work, we propose STGATE (Spatio-Temporal Graph Attention Network with Transformer Encoder), a unified multi-objective model that simultaneously performs next-hour demand forecasting, user behavior clustering, grid stress prediction, and anomaly detection within a single integrated architecture. By modeling charging stations as nodes in a spatial graph and incorporating an explainability framework, the proposed approach not only improves predictive performance but also provides meaningful and actionable insights. The system is further supported by a React-based dashboard for real-time monitoring, interactive visualization, and efficient management of smart grid operations.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

II. LITERATURE REVIEW

A. EV Charging Demand Forecasting

Accurate demand estimation plays a central role in the effective operation of EV charging stations and overall energy management systems. Prior research [1] has explored several forecasting strategies, including probabilistic and machine learning approaches designed to handle uncertainty and variability in demand. Even so, challenges such as limited datasets, irregular usage patterns, and the absence of standardized evaluation metrics continue to affect performance and model generalization. A broader review [2] grouped existing methods into statistical, simulation-driven, and data-centric techniques, suggesting that better results can be achieved by integrating insights across different operational levels and data sources. Similarly, another study [3] proposed a TOPSIS-based framework to guide the selection of forecasting models across varying time horizons and application scenarios.

Time-series techniques remain widely used in this domain due to their ability to model sequential dependencies. Comparative analyses [9] indicate that deep learning models like LSTM tend to outperform classical methods such as ARIMA when sufficient historical data is available and properly preprocessed. Work based on SARIMA models [10] further highlights the importance of accounting for seasonal patterns and periodic trends to improve prediction accuracy. More recent approaches [11] have introduced global models capable of learning shared temporal behavior across multiple stations, thereby improving scalability and consistency. Hybrid architectures, including Transformer-based extensions [5][12], have also demonstrated improved performance by incorporating contextual features such as weather conditions, temporal signals, and calendar-based variations.

B. Ensemble and Machine Learning Approaches

Ensemble techniques have gained attention for their ability to improve prediction accuracy and robustness by combining multiple models with complementary strengths. A weighted ensemble approach [7], integrating algorithms such as Random Forest, LightGBM, XGBoost, and neural networks, achieved strong performance in predicting both energy consumption and session duration across diverse datasets. The use of SHAP analysis in such studies also highlights the growing importance of model interpretability and feature importance analysis in real-world deployments. Other works [13][14] explored hybrid machine learning setups for predicting charging patterns, availability, and waiting times, showing promising and consistent results across multiple evaluation benchmarks. Additionally, Attention-based LSTM models [8] improve short-term demand forecasting with better clarity and flexibility.

C. Graph Neural Network Applications

Graph-based approaches provide a natural and scalable way to model relationships between charging stations, especially in urban environments where spatial dependencies are significant. A data-driven study [4] used tensor decomposition to analyze charging behavior across space and time, identifying congestion zones, underutilized regions, and redundant infrastructure patterns. This reinforces the limitation of treating stations independently and supports the use of graph representations for more holistic modeling. Further research [15][16] shows that combining Graph Neural Networks with temporal models significantly enhances prediction performance by capturing both spatial and temporal correlations simultaneously. In addition, anomaly detection frameworks such as Grid Sentinel [17] have demonstrated effectiveness in identifying abnormal charging activities, faults, and irregular usage patterns within smart grid environments.

D. User Behavior Analysis and Station Utilization

Analyzing user charging behavior is essential for improving infrastructure efficiency, reducing waiting times, and optimizing station placement strategies. Research based on data analytics methodologies like CRISP-DM [18] has provided valuable insights into fleet charging patterns across organizations and large-scale deployments. Similarly, machine learning-based systems [19] have been developed for real-time prediction of station availability, leading to improved recommendation systems, better user experience, and enhanced operational efficiency across evaluation metrics.

E. Smart Grid Analytics and Anomaly Detection

As EV charging systems become more integrated with smart grids, advanced analytics and intelligent control mechanisms are increasingly required to ensure stability and efficiency. A data-driven framework [20] demonstrated how EV charging stations connected to microgrids can actively participate in demand-response strategies during peak periods, thereby reducing grid stress and improving energy utilization. This highlights the need for predictive models that consider both energy demand and grid stability while supporting real-time decision-making.

F. Research Gap and Motivation

Although significant progress has been made, most existing approaches still address individual components such as forecasting, behavior analysis, or anomaly detection in isolation [7][8]. There is a lack of a unified framework that integrates spatial relationships and temporal dynamics while also ensuring interpretability. To bridge this gap, this work introduces STGATE, a unified spatio-temporal architecture with an embedded explainability layer. The proposed model is evaluated using real-world data collected from 462 charging stations across five U.S. cities.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

III. METHODOLOGY

A. Dataset and Preprocessing

This study is based on a dataset containing 1,320 EV charging sessions collected from 462 stations across five major U.S. cities—Houston, San Francisco, Los Angeles, Chicago, and New York—during January and February 2024. Each session record includes details such as station ID, location, energy delivered, session duration, charging rate, initial state of charge, SoC gap, temperature, and user category (Commuter, Casual Driver, or Long-Distance Traveler).

From these records, station-level features were derived and organized into sequences using a sliding window of six time steps. Each sequence captures five key attributes: total energy, session count, average SoC gap, temperature, and charging rate. Before training, all features were normalized using StandardScaler. The dataset was split into training (70%), validation (15%), and testing (15%) subsets.

B. Graph Construction

Charging stations are modeled as nodes within a graph structure, where connections are defined based on geographic closeness and similarity in usage patterns. Each node is described using 13 aggregated features, including energy statistics, session counts, charging behavior, peak usage, user distribution, charger types, and city encoding. The graph is implemented using PyTorch Geometric, allowing efficient handling of node relationships through an edge index representation.

C. System Architecture

Fig. 1 shows the overall STGATE system. The proposed STGATE framework follows a dual-encoder design that combines spatial and temporal learning. A Graph Attention Network (GAT) captures interactions between stations, while a Transformer Encoder processes historical usage sequences. The outputs from both components are merged through a fusion layer and passed to multiple prediction heads responsible for demand estimation, behavior analysis, grid stress evaluation, and anomaly detection.

The model takes two inputs: a graph-based representation of stations and time-series sequences describing past activity. After preprocessing, each input is passed through its respective encoder. The GAT focuses on spatial dependencies, whereas the Transformer learns temporal trends. Their outputs are combined into a shared representation, which is further processed for multi-task predictions. An explainability module is also integrated to provide insights into model behavior.

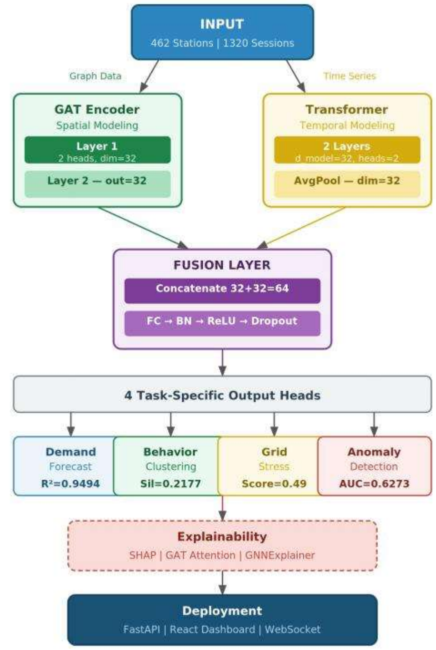

Fig. 1. STGATE system architecture

D. Model Architecture Details

STGATE consists of four main components. The GAT encoder uses two layers with dual attention heads and a hidden size of 32, along with batch normalization, ELU activation, and dropout to improve generalization. The Transformer encoder also contains two layers with a model dimension of 32 and a feedforward size of 128, followed by adaptive average pooling to produce a compact representation.

The outputs of both encoders are concatenated and passed through a fully connected layer to create a unified feature vector. This vector is then used by four task-specific heads: one for predicting next-hour energy demand, one for generating embeddings for user behavior clustering, one for estimating grid stress as a continuous score, and one for classifying anomalies.

E. Training Configuration

The model is implemented using PyTorch and PyTorch Geometric. Training is carried out for up to 100 epochs with a batch size of 8. Early stopping is applied to prevent overfitting, and the Adam optimizer is used with a learning rate of 1e-3 and weight decay of 1e-4. A learning rate scheduler reduces the rate when validation performance plateaus.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

The overall loss function combines multiple objectives: mean squared error for demand and grid predictions, and cross-entropy loss for anomaly detection. Behavior clustering is performed separately using K-Means on the learned embeddings, with the number of clusters selected based on the Silhouette Score.

F. Explainability Framework

To make the system more interpretable, a three-level explainability mechanism is included. SHAP values are used to identify important input features for demand predictions. Node importance scores highlight which stations have the most influence on overall results. Additionally, attention weights from the GAT layers are visualized to show how stations interact with each other. These explanations are accessible through a dedicated API endpoint.

G. Deployment

The trained model is deployed as a FastAPI backend that provides endpoints for all major functionalities, including prediction, clustering, anomaly detection, and explainability. A React-based dashboard connects to this backend using WebSockets and updates data every few seconds, enabling real-time monitoring. The system is hosted online using a Cloudflare tunnel, allowing remote access without requiring a dedicated server.

IV. RESULTS AND DISCUSSION

A. Demand Forecasting Performance

The proposed STGATE model delivers strong performance in next-hour energy demand prediction, achieving an R^2 score of 0.9494 along with an MAE of 4.05 kWh, RMSE of 5.33 kWh, and MAPE of 17.14%. An R^2 value close to 0.95 indicates that the model captures nearly all variability in demand, which highlights the effectiveness of combining spatial and temporal learning. Compared to conventional LSTM-based approaches, which typically report R^2 values in the range of 0.80–0.92 for similar tasks [9], STGATE performs noticeably better. This improvement can be attributed to the inclusion of graph-based spatial modeling, which allows the system to learn dependencies between nearby charging stations—something purely temporal models often miss.

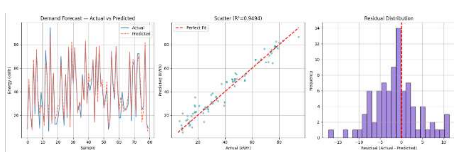

Fig. 2. Demand Forecast — Actual vs Predicted (R^2=0.9494)

B. User Behavior Clustering

The behavior modeling component generates 16-dimensional embeddings that were grouped using K-Means clustering. The optimal number of clusters was found to be two, resulting in a Silhouette Score of 0.2177. While the score is moderate, it still indicates meaningful structure in the data, which is expected given the gradual and overlapping nature of real-world user behavior. The clusters broadly separate users into high-intensity and low-intensity charging profiles. High-intensity users typically include long-distance travelers and those relying on fast charging, whereas low-intensity users are mostly commuters or occasional drivers. These patterns can help guide infrastructure planning—for example, identifying where higher power capacity is needed or where demand-response strategies can be applied.

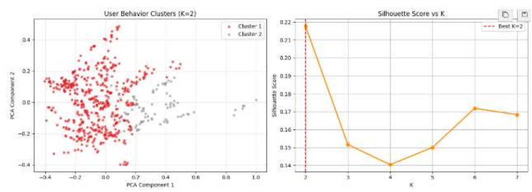

Fig. 3. User Behavior Clusters (K=2) and Silhouette Score vs K

C. Grid Stress Prediction

The grid stress module produced an average stress score of 0.4925 across all stations, suggesting that the overall system operates under moderate load conditions. However, 20 stations were identified as high-stress locations, with scores exceeding 0.7. These stations are primarily associated with areas that have a higher density of fast charging usage. Detecting such hotspots early is useful for grid operators, as it enables proactive measures like redistributing load or implementing load management strategies to maintain system stability.

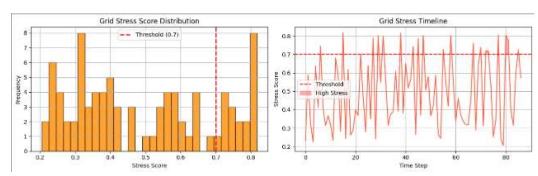

Fig. 4. Grid Stress Score Distribution and Timeline

D. Anomaly Detection

For anomaly detection, the model achieved an AUC-ROC of 0.6273, indicating a reasonable but not highly strong ability to differentiate between normal and abnormal charging sessions. Anomalies were defined based on sessions with unusually high SoC gaps, specifically those above the 90th percentile. The relatively modest performance is partly due to class imbalance, as anomalous events are rare in real-world datasets [8]. To improve this component, future work may incorporate techniques such as oversampling (e.g., SMOTE) or modified loss functions like focal loss to better handle imbalance and improve detection sensitivity.

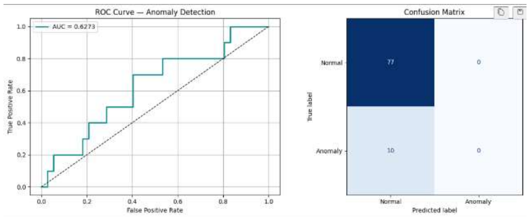

Fig. 5. ROC Curve (AUC=0.6273) and Confusion Matrix

E. Summary of Evaluation Metrics

The STGATE model is evaluated across multiple tasks, including demand forecasting, clustering, grid stress analysis, and anomaly detection, to assess its overall effectiveness. For demand forecasting, the model performs very well, achieving a high R^2 score of 0.9494, which shows that the predicted values closely match actual demand. The error values, including MAE (4.05 kWh) and RMSE (5.33 kWh), are low, indicating accurate predictions, while the MAPE of 17.14% is acceptable given the natural variability in EV charging behavior. In terms of clustering, the Silhouette score of 0.2177 suggests moderate separation between groups, with $k = 2$ providing a reasonable division of user behavior patterns.

For grid stress analysis, the average stress score of 0.4925 indicates a moderate load on the system, and 20 instances are identified as high-stress events, showing the model's ability to highlight potential pressure points in the grid. In anomaly detection, the AUC-ROC score of 0.6273 reflects fair performance in detecting unusual charging activities. While this can be improved further, the overall results demonstrate that STGATE works effectively across different tasks, offering a practical and balanced solution for managing and analyzing EV charging systems in real-world scenarios.

TABLE I
STGATE EVALUATION METRICS

Task	Metric	Value	Result
Demand Forecast	R^2	0.9494	Excellent
Demand Forecast	MAE (kWh)	4.05	Low Error
Demand Forecast	RMSE (kWh)	5.33	Low Error
Demand Forecast	MAPE (%)	17.14	Acceptable
Clustering	Silhouette	0.2177	Moderate
Clustering	k	2	Optimal
Grid Stress	Mean Score	0.4925	Moderate
Grid Stress	High-Stress	20	Flagged
Anomaly Det.	AUC-ROC	0.6273	Fair

V. CONCLUSION

This paper presented STGATE, a spatio-temporal deep learning framework designed to handle multiple aspects of EV charging infrastructure in a unified and efficient manner. The model combines a Graph Attention Network for capturing spatial relationships between stations and a Transformer Encoder for learning complex temporal usage patterns over time. By integrating these components, STGATE is able to perform four key tasks simultaneously: next-hour demand forecasting, user behavior clustering, grid stress prediction, and anomaly detection. Unlike many existing approaches that treat these problems separately, this unified design simplifies the overall system architecture, reduces computational redundancy, improves coordination between tasks, and makes it more practical and scalable for real-world deployment scenarios.

The model was evaluated on 1,320 charging sessions collected from 462 stations across five U.S. cities, providing a diverse and realistic dataset. The results show strong performance, particularly in demand forecasting with an R^2 of 0.9494, indicating high predictive accuracy and reliable generalization within the dataset. The clustering output revealed two broad user groups—high-intensity and low-intensity users—which can be useful for planning infrastructure capacity, optimizing resource allocation, and designing demand-response strategies. The grid stress analysis also helped identify high-load stations and potential bottlenecks, enabling better operational planning, while anomaly detection showed moderate but acceptable performance given the inherent noise and imbalance in the data.

An important aspect of STGATE is its built-in explainability, which enhances its practical usability and trustworthiness. By combining SHAP analysis, attention visualization, and node importance scoring, the model provides insights that are easier to interpret compared to typical black-box approaches. This makes it more useful for planners and operators who need to understand the reasoning behind predictions and make informed decisions in dynamic environments. The system is further supported by a FastAPI backend and a React-based dashboard, enabling real-time monitoring, interactive visualization, and seamless integration into existing smart grid systems and operational workflows.

That said, there are still some limitations that need to be addressed. The dataset is relatively small and limited to a short time period, which may affect the model's ability to generalize across different regions and long-term conditions. The anomaly detection performance is also impacted by class imbalance, and the use of a static graph limits the model's ability to capture evolving relationships between stations over time. Future work will focus on expanding the dataset, improving anomaly detection using techniques like oversampling and advanced loss functions, and exploring dynamic graph-based approaches to better model real-world variability, seasonal trends, and system evolution.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 04 , April - 2026

REFERENCES

[1] M. Rashid, T. Elfouly, and N. Chen, "A Comprehensive Survey of Electric Vehicle Charging Demand Forecasting Techniques," *IEEE Open J. Veh. Technol.*, vol. 5, pp. 1348–1373, 2024, doi: 10.1109/OJVT.2024.3457499.

[2] X. Lin *et al.*, "Electric Vehicle Charging Load Modeling: A Survey, Trends, Challenges and Opportunities," Oct. 30, 2025, arXiv:2511.03741. doi: 10.48550/arXiv.2511.03741.

[3] M. Kim, W. Y. Park, H. S. Moon, and Y.-S. Kim, "A Review of Electric Vehicle Charging Demand Forecasting Models by Prediction Horizon: A Multi-Criteria Decision Analysis Approach," *J. Electr. Eng. Technol.*, Jan. 2026, doi: 10.1007/s42835-025-02556-0.

[4] Y. Yang, Y. Zhang, and X. Meng, "A Data-Driven Approach for Optimizing the EV Charging Stations Network," *IEEE Access*, vol. 8, pp. 118572–118592, 2020, doi: 10.1109/ACCESS.2020.3004715.

[5] A. Hussain, V. Eswarakrishnan, A. Aslam, and S. Tripura, "Charging stations demand forecasting using LSTM based hybrid transformer model," *Sci. Rep.*, vol. 15, no. 1, p. 36639, Oct. 2025, doi: 10.1038/s41598-025-20421-y.

[6] R. N. Akshayaa, S. S A, and R. R, "Session and energy forecasting for electric vehicle charging station using custom weighted ensemble machine learning," *Results Eng.*, vol. 26, p. 105104, Jun. 2025, doi: 10.1016/j.rineng.2025.105104.

[7] S. Sanami, H. Mosalli, Y. Yang, H.-G. Yeh, and A. G. Aghdam, "Demand Forecasting for Electric Vehicle Charging Stations using Multivariate Time-Series Analysis," Feb. 22, 2025, *arXiv*: arXiv:2502.16365. doi: 10.48550/arXiv.2502.16365.

[8] V. T. Kesavan, Md. J. Hossen, R. Gopi, and E. R. Joseph, "Anomaly detection with grid sentinel framework for electric vehicle charging stations in a smart grid environment," *Sci. Rep.*, vol. 15, no. 1, p. 15774, May 2025, doi: 10.1038/s41598-025-00400-z.

[9] Y. Kim and S. Kim, "Forecasting Charging Demand of Electric Vehicles Using Time-Series Models," *Energies*, vol. 14, no. 5, p. 1487, Mar. 2021, doi: 10.3390/en14051487.

[10] K. C. Akshay, G. H. Grace, K. Gunasekaran, and R. Samikannu, "Power consumption prediction for electric vehicle charging stations and forecasting income," *Sci. Rep.*, vol. 14, no. 1, p. 6497, Mar. 2024, doi: 10.1038/s41598-024-56507-2.

[11] V. Kaur, D. Sharma, S. K. Dwivedi, M. K. Yadav, A. K. L. Singh, and P. Sharma, "Prediction of Electric Vehicle Charging Using Machine Learning," *SSRN Electron. J.*, 2024, doi: 10.2139/ssrn.4495876.

[12] T. Van Etten, V. Degeler, and D. Luo, "Large-Scale Forecasting of Electric Vehicle Charging Demand Using Global Time Series Modeling:," in *Proceedings of the 10th International Conference on Vehicle Technology and Intelligent Transport Systems*, Angers, France: SCITEPRESS - Science and Technology Publications, 2024, pp. 40–51. doi: 10.5220/0012555400003702.

[13] Z. Yi, X. C. Liu, R. Wei, X. Chen, and J. Dai, "Electric vehicle charging demand forecasting using deep learning model," *J. Intell. Transp. Syst.*, vol. 26, no. 6, pp. 690–703, Nov. 2022, doi: 10.1080/15472450.2021.1966627.

[14] A. Hussain, Q.-C. Lu, S. S. Rizvi, S. Wang, and S. J. Kwon, "Short term demand forecasting of electric vehicle charging stations using context aware temporal transformer model," *Sci. Rep.*, vol. 15, no. 1, p. 36652, Oct. 2025, doi: 10.1038/s41598-025-20557-x.

[15] F. Gonçalves, L. De Abreu Borges, and R. Batista, "Electric Vehicle Charging Data Analytics of Corporate Fleets," *World Electr. Veh. J.*, vol. 13, no. 12, p. 237, Dec. 2022, doi: 10.3390/wevj13120237.

[16] M. Manai, B. Sellami, and S. Ben Yahia, "Accurate Recommendation of EV Charging Stations Driven by Availability Status Prediction:," in *Proceedings of the 19th International Conference on Software Technologies*, Dijon, France: SCITEPRESS - Science and Technology Publications, 2024, pp. 351–358. doi: 10.5220/0012752600003753.

[17] A. Hafeez, R. Alammari, and A. Iqbal, "Utilization of EV Charging Station in Demand Side Management Using Deep Learning Method," *IEEE Access*, vol. 11, pp. 8747–8760, 2023, doi: 10.1109/ACCESS.2023.3238667.

[18] A. Kumari, "Machine Learning-Based Prediction for Ev Charging Station Availability and Wait-Time Estimation," 2024, *SSRN*. doi: 10.2139/ssrn.5035693.

[19] Z. Tang, Y. Cui, Q. Hu, M. Liu, W. Rao, and X. Liu, "Electric vehicle charging load prediction based on graph attention networks and autoformer," *J. Eng.*, vol. 2024, no. 9, p. e70009, Sep. 2024, doi: 10.1049/tje2.70009.

[20] 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," *J. Intell. Fuzzy Syst.*, vol. 46, no. 1, pp. 821–836, Jan. 2024, doi: 10.3233/JIFS-231775.