International Journal of Creative and Open Research in Engineering and Management
ISSN: **3108-1754 (Online)**
Volume 02 Issue 05 May-2026 | Impact Factor: 3.5

A Framework for Policy-Driven Integration of Renewable Energy in Electric Vehicles

1st Suman
Department of Electrical and Electronics Engineering
Maharaja Surajmal Institute of Technology (MSIT), New Delhi, India

How to Cite this Article:
Suman, (2026). A Framework for Policy-Driven Integration of Renewable Energy in Electric Vehicles. International Journal of Creative and Open Research in Engineering and Management, <i>02</i>(05).
https://doi.org/10.55041/ijcope.v2i5.542

https://doi.org/10.55041/ijcope.v2i5.542

Abstract—The accelerating deployment of electric vehicles (EVs) alongside the rapid expansion of renewable energy sources (RES) presents both transformative opportunities and complex operational challenges for modern power systems. Existing policy mechanisms, such as time-of-use (ToU) tariffs, renewable portfolio standards (RPS), net metering credits, and vehicle-to-grid (V2G) incentives, are typically designed and evaluated in isolation from real-time grid dynamics, leaving a critical gap between regulatory intent and operational effectiveness. This paper proposes a novel deep learning-based analytical framework—designated PolicyRE-EV—designed to bridge this gap by coupling renewable energy availability forecasting with policy-sensitive EV charging demand modeling. The proposed framework integrates a Transformer-based multivariate forecasting backbone with a dedicated Policy Encoding Module that represents regulatory parameters as differentiable soft constraints within the training objective. Three operational policy scenarios are investigated: (1) unregulated baseline EV charging, (2) ToU tariff-regulated smart charging, and (3) full renewable-aligned V2G dispatch. Experimental benchmarks against Long Short-Term Memory (LSTM), Bidirectional LSTM (BiLSTM), and Random Forest (RF) baseline models demonstrate that PolicyRE-EV achieves superior performance across Mean Absolute Error (MAE), Root Mean Square Error (RMSE), Mean Absolute Percentage Error (MAPE), Renewable Utilization Rate (RUR), and CO_2 Displacement Index (CDI). The results confirm that embedding policy parameters within deep learning architectures significantly enhances alignment between EV charging patterns and renewable generation windows, offering a scalable and policy-compliant tool for distribution network operators and energy planners.

Index Terms—Renewable Energy Integration, Electric Vehicle Charging, Policy Framework, Transformer Architecture, Vehicle-to-Grid (V2G), Time-of-Use Tariffs, Smart Grid, Deep Learning Forecasting, Demand-Side Management.

I. INTRODUCTION

The global energy transition is fundamentally reshaping the operational landscape of electrical power systems. The concurrent rise of renewable energy sources—particularly solar photovoltaics (PV) and wind power—and the accelerating electrification of the transportation sector through electric vehicles (EVs) are two defining trends of this transformation. According to the International Energy Agency, global EV stock surpassed 40 million units by 2023, with projections indicating over 240 million EVs on the road by 2030 under stated policy scenarios [1]. Simultaneously, renewable energy accounted for nearly 30% of global electricity generation in 2023, a share expected to grow substantially under net-zero emission pathways.

While the convergence of EV adoption and renewable energy deployment holds immense promise for decarbonization, it introduces significant technical and regulatory challenges. EV charging is inherently stochastic, driven by driver behavior, travel patterns, and ambient conditions. Renewable energy generation is similarly intermittent, subject to solar irradiance variability and wind fluctuations. When these two volatile elements interact within a distribution network lacking intelligent coordination, the result is compounding uncertainty, potential grid instability, and suboptimal renewable utilization.

Policy frameworks serve as the primary instrument through which governments and regulators seek to align the behavior of EV users with grid operational objectives and renewable energy goals. Time-of-use (ToU) tariffs incentivize off-peak charging, renewable portfolio standards (RPS) mandate minimum renewable generation shares, net metering schemes reward distributed solar prosumers, and vehicle-to-grid (V2G) programs enable EVs to discharge stored energy back to the grid during demand peaks. However, the design and evaluation of these policies has historically relied on static optimization models or simplified simulation frameworks that fail to capture the nonlinear temporal complexity of real-world EV and renewable energy interactions [2].

Deep learning architectures, particularly Transformer-based models utilizing self-attention mechanisms, have demonstrated exceptional performance in modeling complex multivariate time-series data [3]. Their ability to capture long-range temporal dependencies and feature interactions makes them uniquely suited to simultaneously represent renewable generation profiles, EV demand patterns, and the time-varying effects of regulatory interventions. This paper exploits these properties to propose PolicyRE-EV, a framework that treats policy parameters as first-class inputs within a Transformer-based forecasting architecture, enabling the model to learn how regulatory constraints reshape charging behavior and renewable consumption.

The primary contributions of this paper are as follows:

- A novel deep learning framework, PolicyRE-EV, that jointly forecasts renewable energy generation and policy-sensitive EV charging demand using a Transformer backbone with patch-based temporal encoding.
- A Policy Encoding Module (PEM) that transforms regulatory parameters—ToU tariffs, V2G incentives, RPS obligations, and net metering credits—into differentiable soft constraints embedded within the model's composite learning objective.
- Systematic empirical evaluation across three policy scenarios using a comprehensive dataset combining solar generation records, EV charging session logs, and regulatory filing data.
- Introduction of two domain-specific evaluation metrics—Renewable Utilization Rate (RUR) and CO_2 Displacement Index (CDI)—to quantify the practical environmental benefits of policy-aligned EV charging.

II. LITERATURE REVIEW

Research at the intersection of renewable energy integration and electric vehicle management has grown considerably over the past decade, spanning disciplines from power systems engineering and transportation science to machine learning and public policy. Early approaches to coordinating EV charging with renewable generation employed deterministic optimization methods such as linear programming and mixed-integer programming, which produced optimal charging schedules under simplifying assumptions about user behavior and generation certainty [4]. While these methods offered interpretable solutions, their computational intractability at scale and inability to handle real-time stochasticity limited practical deployment.

Stochastic optimization and robust programming methods were subsequently introduced to account for uncertainty in both renewable generation and EV availability. Scenario-based approaches using Monte Carlo sampling have been used to evaluate policy interventions such as ToU tariff designs and V2G incentive structures under probabilistic demand and generation forecasts [5]. However, these methods typically require the specification of parametric probability distributions, imposing restrictive assumptions that may not reflect empirical data distributions, particularly in high-penetration EV environments.

The application of machine learning to renewable energy forecasting has a well-established history, with Artificial Neural Networks (ANNs), Support Vector Regression (SVR), and ensemble methods such as Random Forest demonstrating competitive accuracy for solar and wind power prediction [6]. Long Short-Term Memory (LSTM) networks, introduced by Hochreiter and Schmidhuber [7], became the dominant deep learning approach for energy time-series forecasting due to their ability to model sequential dependencies through gated memory cells. Bidirectional LSTM (BiLSTM) variants, which process sequences in both forward and backward temporal directions, have shown further improvements in capturing asymmetric temporal patterns relevant to renewable generation and load demand [8].

Transformer architectures, first proposed for natural language processing by Vaswani et al. [9], have been adapted with increasing success for time-series forecasting. The Informer model [10] introduced sparse self-attention to handle long input sequences efficiently, while the PatchTST model [11] proposed segmenting time series into non-overlapping patches before applying self-attention, achieving state-of-the-art performance on multiple long-horizon energy forecasting benchmarks. The Temporal Fusion Transformer (TFT) [12] further extended this paradigm by incorporating multi-horizon forecasting with interpretable attention weights and explicit support for static and time-varying covariates.

Despite these advances in forecasting methodology, the explicit integration of policy parameters into deep learning models remains limited. Most existing studies treat policy scenarios as external simulation conditions evaluated through separate model runs, rather than embedding regulatory parameters as active inputs that the model can learn to respond to. This design choice prevents the model from capturing the nuanced behavioral adaptations—changes in charging timing, session duration, and grid injection patterns—that policy instruments induce in EV users and grid operators. A small number of recent studies have begun exploring reinforcement learning approaches for adaptive EV charging policy optimization [13], but these methods require extensive environment simulation and do not readily incorporate existing policy documents or regulatory schedules as structured inputs. The present work addresses this gap through a supervised deep learning approach that encodes policy parameters as structured features within a Transformer architecture.

III. PROBLEM FORMULATION

A. System State Representation

The joint system state at time step t is represented by the multivariate feature vector:

$$S_t = [G_t, L_t^{EV}, L_t^{base}, P_t^{ToU}, I_t^{V2G}, R_t^{RPS}, N_t^{NEM}, T_t, H_t, D_t, W_t] \qquad (1)$$

where G_t is the renewable generation output (kW) at time t; L_t^{EV} and L_t^{base} denote EV charging load and baseline residential demand (kW) respectively; P_t^{ToU} is the prevailing time-of-use electricity price (currency/kWh); I_t^{V2G} is a binary indicator of active V2G incentive; R_t^{RPS} represents the renewable portfolio standard compliance level; N_t^{NEM} is the net metering credit (currency/kWh); and T_t, H_t, D_t, W_t encode temperature, humidity, day-of-week, and weekend/holiday flag respectively.

B. Forecasting Problem Statement

Given a lookback window of length n, the framework learns a nonlinear mapping f parameterized by θ to predict the joint output at horizon h:

$$\hat{S}_{t+h} = f(S_{t-n+1:t}; \theta) \qquad (2)$$

The output vector $\hat{S}_{t+h} = [\hat{G}_{t+h}, \hat{L}_{t+h}^{EV}]$ contains forecasted renewable generation and EV charging demand. The model parameters are optimized by minimizing a composite loss function:

$$L_{total} = L_{forecast} + \lambda_1 \cdot L_{policy} + \lambda_2 \cdot L_{V2G} \qquad (3)$$

C. Loss Function Components

The forecasting loss measures prediction error over N training samples:

$$L_{forecast} = (1/N)\ \Sigma_i (S_i - \hat{S}_i)^2 \qquad (4)$$

The policy constraint loss penalizes EV charging scheduled during high-tariff periods when renewable generation is insufficient to cover demand:

$$L_{policy} = (1/N) \sum_i \max(0, \hat{L}_{i,EV} \cdot P_i^{ToU} - \alpha \cdot G_i) \quad (5)$$

where α is a scalar representing the renewable absorption coefficient of the feeder. The V2G incentive loss promotes battery discharge during high-tariff, low-generation intervals:

$$L_{V2G} = (1/N) \sum_i I_i^{V2G} \cdot \max(0, E_i^{bat} - E_i^{discharge}) \quad (6)$$

where E_i^{bat} is the available battery state-of-charge and $E_i^{discharge}$ is the target energy injection quantity.

D. Renewable Utilization Rate (RUR)

To assess renewable energy consumption effectiveness, the Renewable Utilization Rate is defined as the fraction of total renewable generation absorbed by EV charging:

$$RUR = [\sum_t \min(L_t^{EV}, G_t)] / [\sum_t G_t] \times 100\% \quad (7)$$

E. CO₂ Displacement Index (CDI)

The carbon displacement attributable to renewable-aligned EV charging, normalized by fleet capacity, is:

$$CDI = \sum_t [L_t^{EV} \cdot (EF_{grid} - EF_{RE})] / (T \cdot C_{fleet}) \quad (8)$$

where EF_{grid} and EF_{RE} are the grid-average and renewable energy emission factors (kg CO_2/kWh), T is the evaluation period, and C_{fleet} is the aggregate fleet energy capacity (kWh).

IV. PROPOSED POLICYREEV FRAMEWORK ARCHITECTURE

A. Overall Design

PolicyRE-EV is structured as a five-module architecture: (1) Patch-Based Input Embedding, (2) Policy Encoding Module (PEM), (3) Temporal Transformer Encoder, (4) Cross-Attention Policy Fusion Layer, and (5) Composite Regression Output Head. This design enables parallel processing of temporal dynamics and regulatory context, which are subsequently fused through cross-attention before generating joint forecasts of renewable generation and EV demand.

The overall data flow can be summarized as: Multivariate System State Input → Patch Embedding + Positional Encoding → Temporal Transformer Encoder → Cross-Attention Fusion (with Policy Encoding Module) → Regression Output Head → Joint Renewable and EV Demand Forecast.

B. Patch-Based Input Embedding

The input sequence of length n is divided into non-overlapping patches of length p, inspired by the PatchTST paradigm. Each patch is linearly projected into a d-dimensional embedding space:

$$Z^{patch}_k = S_{t-n+kp:t-n+(k+1)p} \cdot W_e, \quad k = 0, 1, \ldots, (n/p)-1 \quad (9)$$

Sinusoidal positional encodings are added to preserve temporal ordering:

$$PE(pos, 2i) = \sin(pos / 10000^{2i/d}) \quad (10)$$

$$PE(pos, 2i+1) = \cos(pos / 10000^{2i/d}) \quad (11)$$

C. Policy Encoding Module (PEM)

Policy parameters are processed through a dedicated lightweight encoder. The policy vector at time t is defined as:

$$P_t = [P_t^{ToU}, I_t^{V2G}, R_t^{RPS}, N_t^{NEM}] \quad (12)$$

These regulatory inputs are projected into the shared embedding space via a two-layer feed-forward network with ReLU activation:

$$P^{enc} = \mathrm{ReLU}(\mathrm{ReLU}(P_t \cdot W_{p1} + b_{p1}) \cdot W_{p2} + b_{p2}) \quad (13)$$

This projection ensures that policy representations are dimensionally compatible with temporal embeddings for downstream cross-attention computation.

D. Temporal Transformer Encoder

The temporal branch consists of L stacked Transformer encoder layers, each implementing multi-head scaled dot-product self-attention:

$$\mathrm{Attention}(Q, K, V) = \mathrm{Softmax}(QK^T / \sqrt{d_k})\, V \qquad (14)$$

$$\mathrm{MultiHead}(Q, K, V) = \mathrm{Concat}(head_1, \ldots, head_H) \cdot W_O \qquad (15)$$

Each attention head applies independent learned projection matrices $W_i^Q, W_i^K, W_i^V \in \mathbb{R}^{d \times d_k}$. A position-wise feed-forward network follows each attention sub-layer:

$$\mathrm{FFN}(x) = \max(0, x \cdot W_1 + b_1) \cdot W_2 + b_2 \qquad (16)$$

Residual connections and layer normalization are applied after each sub-layer to stabilize gradient flow during training.

E. Cross-Attention Policy Fusion Layer

The temporal encoder output H_{temp} serves as the query, while the policy encoder output P^{enc} provides keys and values in a cross-attention operation that enables the model to selectively attend to relevant policy signals:

$$H_{fused} = \mathrm{Softmax}(H_{temp} \cdot (P^{enc})^T / \sqrt{d_k}) \cdot P^{enc} \qquad (17)$$

This formulation allows the model to dynamically weight different policy parameters according to their temporal relevance—for example, assigning higher attention weights to V2G incentive signals during periods of high solar generation.

F. Composite Regression Output

The fused representation H_{fused} passes through a fully connected layer to produce the joint prediction output:

$$\hat{S}_{t+h} = [\hat{G}_{t+h}, \hat{L}_{t+h}^{EV}] = W_o \cdot H_{fused} + b_o \qquad (18)$$

The composite loss defined in Equation (3) is minimized using the Adam optimizer with weight decay regularization and cosine annealing learning rate scheduling.

V. DATA COLLECTION AND PREPROCESSING

A. Dataset Description

The experimental dataset integrates three complementary data sources. Solar PV generation time series at 15-minute resolution are sourced from the National Renewable Energy Laboratory (NREL) Solar Radiation Database (NSRDB) for a representative mid-latitude urban distribution feeder over a 24-month period (January 2022–December 2023). EV charging demand profiles are constructed by combining travel behavior microdata from the National Household Travel Survey (NHTS) with charging session records conforming to the Open Charge Point Protocol (OCPP), representing a fleet of 500 mixed-class EVs (sedans, SUVs, light commercial) connected to a residential low-voltage feeder. Policy parameter time series—comprising ToU tariff schedules, V2G incentive activation windows, and RPS compliance levels—are derived from publicly available regulatory filings of a representative state public utility commission.

B. Data Cleaning Procedures

Raw multi-source data requires careful preprocessing prior to model ingestion. The following sequential pipeline is applied:

- Missing Value Treatment: Gaps spanning fewer than four consecutive 15-minute intervals are filled via linear interpolation. Longer gaps are addressed using seasonal decomposition-based reconstruction that extracts and reapplies weekly periodic patterns from adjacent clean data segments.

- Outlier Detection and Correction: The Interquartile Range (IQR) method identifies anomalous values lying beyond $Q1 - 1.5 \cdot IQR$ or $Q3 + 1.5 \cdot IQR$. Detected outliers are replaced with locally computed exponentially weighted moving averages to preserve temporal continuity.
- Temporal Alignment: All heterogeneous data streams are resampled to a unified 15-minute resolution. Discrete policy variables are forward-filled, while continuous measurements undergo mean aggregation across sub-periods.

C. Feature Normalization

All continuous input variables are scaled to the unit interval [0, 1] using Min-Max normalization to prevent any single feature from dominating the learning process due to scale differences:

$$X\_norm = (X - X\_min) / (X\_max - X\_min) \qquad (19)$$

Binary indicators (I_t^V2G, weekend flag W_t) retain their original {0, 1} encoding. ToU tariff values are normalized independently relative to their own range to preserve the economic signal structure.

D. Feature Engineering

Beyond raw input variables, the following derived features are constructed to enhance the model's capacity to capture temporal patterns:

- Cyclic Time Encoding: Hour-of-day and month-of-year are transformed into sine-cosine pairs to represent periodic continuity without artificial boundary discontinuities at transitions such as 23:45 to 00:00: $h\_sin = \sin(2\pi \cdot hour/24)$, $h\_cos = \cos(2\pi \cdot hour/24)$.
- Lag Features: Historical values of load and generation at lags of 1 hour (k=4 intervals), 24 hours (k=96), and 7 days (k=672) are included to capture intra-day, daily, and weekly seasonal dependencies.
- Policy Interaction Terms: Multiplicative products of ToU tariff and renewable generation availability are computed to explicitly represent the economic incentive for demand shifting.
- Rolling Statistics: One-hour, six-hour, and 24-hour rolling mean and standard deviation of load and generation are appended to encode local trend and volatility information.

E. Dataset Partitioning and Sequence Construction

The dataset is partitioned chronologically: the first 18 months (75%) form the training set, the subsequent 3 months (12.5%) constitute the validation set for hyperparameter tuning, and the final 3 months (12.5%) serve as the held-out test set. Sliding window sampling with a lookback of $n = 96$ steps (24 hours at 15-minute resolution) and a prediction horizon of $h = 4$ steps (1 hour ahead) is applied to construct input-output sample pairs.

VI. EXPERIMENTAL SETUP

All models are implemented in Python 3.10 using the PyTorch 2.1 deep learning framework. Experiments are conducted on an NVIDIA A100 GPU (40 GB HBM2e memory). The PolicyRE-EV hyperparameter configuration is as follows: embedding dimension $d = 128$, number of attention heads $H = 8$, Transformer encoder layers $L = 4$, patch length $p = 12$ (3-hour segments), dropout rate = 0.1, and policy constraint weights $\lambda_1 = 0.30$, $\lambda_2 = 0.20$. Training employs the Adam optimizer with initial learning rate 1×10^{-3}, cosine annealing decay, batch size 64, and early stopping with patience 20 evaluated on validation loss. All experiments are repeated across five independent random seeds; results report mean ± standard deviation.

Baseline models for benchmarking are: (1) Random Forest with 200 estimators and maximum depth 15; (2) LSTM with two stacked layers of 256 hidden units; and (3) BiLSTM with two layers of 256 units per direction. All baselines receive the same feature set as PolicyRE-EV; policy constraint loss components are exclusive to the proposed architecture.

Three policy scenarios are evaluated:

- Scenario 1 (S1) – Unregulated Baseline: No ToU pricing or V2G incentive is active. EV charging follows unmanaged demand behavior without policy intervention.

- Scenario 2 (S2) – ToU Tariff Regulation: Time-of-use pricing applies with peak tariff from 17:00–21:00 at 3× the off-peak rate, creating a cost incentive to shift charging to solar generation hours.
- Scenario 3 (S3) – Full Renewable-V2G Integration: ToU tariffs are active alongside V2G incentive windows aligned with peak solar generation (10:00–15:00) and morning valley periods, permitting bidirectional energy flows.

VII. RESULTS AND PERFORMANCE METRICS

A. Evaluation Metrics

Forecasting accuracy is quantified using three standard regression error metrics. Mean Absolute Error (MAE) measures the average magnitude of prediction deviations:

$$MAE = (1/N)\ \Sigma_i\ |L\_i - \hat{L}\_i| \qquad (20)$$

Root Mean Square Error (RMSE) applies greater penalty to large prediction errors:

$$RMSE = \sqrt{[(1/N)\ \Sigma_i\ (L\_i - \hat{L}\_i)^2]} \qquad (21)$$

Mean Absolute Percentage Error (MAPE) expresses error relative to actual values:

$$MAPE = (100/N)\ \Sigma_i\ |L\_i - \hat{L}\_i|\ /\ |L\_i| \qquad (22)$$

Renewable integration effectiveness is assessed using RUR (Eq. 7) and CDI (Eq. 8).

B. Comparative Performance Results

Table I presents the full performance comparison across all models and scenarios on the held-out test set.

TABLE I

COMPARATIVE PERFORMANCE OF ALL MODELS ACROSS THREE POLICY SCENARIOS

Model	**Scenario**	**MAE (kW)**	**RMSE (kW)**	**MAPE (%)**	**RUR (%)**	**CDI (kg/kWh)**	**R^2**
Random Forest	S1	0.144	0.191	7.96	40.8	0.079	0.868
Random Forest	S2	0.138	0.183	7.60	48.1	0.094	0.876
Random Forest	S3	0.133	0.176	7.28	52.7	0.110	0.881
LSTM	S1	0.119	0.156	6.47	44.3	0.087	0.899
LSTM	S2	0.110	0.145	5.99	53.0	0.105	0.911
LSTM	S3	0.102	0.136	5.55	58.8	0.122	0.920
BiLSTM	S1	0.108	0.143	5.93	45.9	0.092	0.910
BiLSTM	S2	0.099	0.131	5.44	55.4	0.111	0.924
BiLSTM	S3	0.092	0.122	5.06	62.1	0.129	0.932
PolicyRE-EV (Proposed)	**S1**	**0.088**	**0.116**	**4.71**	**48.6**	**0.096**	**0.943**
PolicyRE-EV (Proposed)	**S2**	**0.076**	**0.101**	**4.18**	**61.8**	**0.122**	**0.956**
PolicyRE-EV (Proposed)	**S3**	**0.063**	**0.084**	**3.44**	**75.2**	**0.158**	**0.969**

Table I: Performance comparison across all models under three policy scenarios (S1 = Unregulated, S2 = ToU Tariff, S3 = Full Renewable-V2G).

Fig. 1. PolicyRE-EV framework architecture: Patch-Based Input Embedding and Policy Encoding Module feed a Temporal Transformer Encoder and Cross-Attention Policy Fusion Layer, producing joint renewable and EV demand forecasts.

Fig. 2. MAE and RMSE comparison bar chart for all models under Scenario 3 (Full Renewable-V2G Integration).

Fig. 3. Renewable Utilization Rate (RUR) across three policy scenarios for all evaluated models, illustrating progressive improvement with policy activation depth.

Fig. 4. 24-hour overlay of actual vs. predicted EV charging demand and solar generation under Scenario 3, demonstrating PolicyRE-EV's alignment between charging peaks and renewable generation windows.

C. Analysis and Discussion

The results in Table I reveal several significant findings. Across all evaluated scenarios, PolicyRE-EV achieves the lowest error values on MAE, RMSE, and MAPE, with the most substantial performance advantages observed in Scenario 3. Compared to the strongest baseline (BiLSTM), the proposed framework reduces MAPE by 32.0% and RMSE by 31.1% under full renewable-V2G integration. This improvement stems from two architectural innovations working in concert: patch-based temporal encoding captures multi-scale periodicity without the gradient vanishing issues inherent in recurrent processing, while the Cross-Attention Policy Fusion Layer enables the model to dynamically attend to relevant policy signals in response to evolving grid conditions.

The RUR results demonstrate that policy design has a material and measurable impact on renewable energy consumption efficiency. PolicyRE-EV achieves a RUR of 75.2% under Scenario 3, representing a 54.7% relative improvement over the unregulated Scenario 1 baseline of 48.6%. This gain substantially exceeds the improvements observed for all baseline models under the same scenario transition, confirming that embedding policy constraints within the learning objective—rather than applying them post-hoc—produces qualitatively superior renewable alignment. The practical implication is that regulatory authorities can leverage policy-aware AI frameworks to design tariff structures and incentive schedules that produce predictable, large-scale shifts in EV charging behavior toward renewable generation windows.

The CO_2 Displacement Index results provide a tangible measure of the environmental benefit achievable through policy-aligned smart charging. Under Scenario 3, PolicyRE-EV achieves a CDI of 0.158 kg CO_2/kWh, equivalent to eliminating approximately 28-30% of the carbon intensity of conventional grid electricity for EV charging. Extrapolating this outcome to a city-scale EV fleet of 500,000 vehicles with average daily consumption of 12 kWh per vehicle yields an estimated annual carbon displacement of approximately 346,000 tonnes CO_2—a reduction equivalent to removing over 75,000 internal combustion engine vehicles from the road annually.

VIII. CONCLUSION

This paper has presented PolicyRE-EV, a deep learning framework designed to advance the policy-driven integration of renewable energy in electric vehicle charging infrastructure. By embedding regulatory parameters—time-of-use tariffs, vehicle-to-grid incentives, renewable portfolio standards, and net metering credits—as differentiable soft constraints within a Transformer-based forecasting architecture, the proposed framework bridges the gap between policy design and operational implementation that has characterized existing approaches.

Comprehensive experimental evaluation across three policy scenarios confirms that PolicyRE-EV consistently outperforms LSTM, BiLSTM, and Random Forest baselines across all evaluated metrics. The full renewable-V2G integration scenario achieves a Renewable Utilization Rate of 75.2% and a CO_2 Displacement Index of 0.158 kg/kWh, representing substantial and scalable environmental benefits. The cross-attention policy fusion mechanism is identified as a key driver of performance, enabling the model to selectively respond to regulatory signals in a temporally adaptive manner.

These results establish that policy-aware deep learning architectures represent a viable and effective pathway for aligning EV charging ecosystems with renewable energy goals, offering distribution system operators, grid planners, and energy regulators a powerful data-driven tool for evidence-based policy design and real-time operational management.

IX. FUTURE SCOPE

The current framework opens several avenues for future investigation. Federated learning extensions would enable multi-feeder deployment without requiring the centralization of sensitive user consumption data, preserving privacy while scaling the framework to regional grid operators. The incorporation of battery state-of-health (SoH) degradation models and stochastic driver availability constraints into the optimization objective would enhance physical realism and improve practical V2G dispatch recommendations. Reinforcement learning-based extensions could transform PolicyRE-EV from a supervised forecasting framework into an adaptive online control system capable of dynamically revising charging schedules in response to real-time renewable availability signals and spot market prices. Furthermore, multi-objective optimization formulations that simultaneously address renewable utilization, voltage profile maintenance, feeder thermal loading, and charging equity across socioeconomic groups represent a natural evolution of the proposed approach. Finally, integrating carbon market price signals and cross-sector coupling with district heating networks and green hydrogen production infrastructure would position the framework within broader whole-energy-system optimization contexts aligned with long-term net-zero emission transition pathways.

REFERENCES

[1] International Energy Agency (IEA), “Global EV Outlook 2024,” IEA Publications, Paris, France, 2024.

[2] P. Morais, H. Kagan, and M. Madrigal, “Policy frameworks for electric vehicle smart charging integration: A critical review,” Renewable and Sustainable Energy Reviews, vol. 178, p. 113265, 2023.

[3] Y. Liu, H. Wu, J. Wang, and M. Long, “Non-stationary Transformers: Exploring the stationarity in time series forecasting,” Advances in Neural Information Processing Systems (NeurIPS), vol. 35, 2022.

[4] S. Agarwal, A. Saxena, and R. Gupta, “Optimal scheduling of EV charging with renewable energy integration using mixed-integer linear programming,” IEEE Transactions on Smart Grid, vol. 13, no. 4, pp. 3042–3055, 2022.

[5] T. Hong and S. Fan, “Probabilistic electric load forecasting: A tutorial review,” International Journal of Forecasting, vol. 32, no. 3, pp. 914–938, 2016.

[6] I. K. Nti, M. Teimeh, and O. Nyarko-Boateng, “Electricity load forecasting: A systematic review,” Journal of Electrical Systems and Information Technology, vol. 7, no. 13, 2020.

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

[8] X. Xiong, Z. Huang, Y. Chen, and J. Sun, “Load forecasting for commercial buildings using BiLSTM–Transformer network,” Symmetry, vol. 16, no. 12, 2024.

[9] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” Advances in Neural Information Processing Systems (NeurIPS), vol. 30, 2017.

[10] H. Zhou, S. Zhang, J. Peng, S. Zhang, J. Li, H. Xiong, and W. Zhang, “Informer: Beyond efficient Transformer for long sequence time-series forecasting,” Proceedings of AAAI Conference on Artificial Intelligence, vol. 35, no. 12, pp. 11106–11115, 2021.

[11] Y. Nie, N. H. Nguyen, P. Sinthong, and J. Kalagnanam, “A time series is worth 64 words: Long-term forecasting with Transformers,” International Conference on Learning Representations (ICLR), 2023.

[12] B. N. Lim, S. Ö. Arik, N. Loeff, and T. Pfister, “Temporal fusion transformers for interpretable multi-horizon time series forecasting,” International Journal of Forecasting, vol. 37, no. 4, pp. 1748–1764, 2021.

[13] M. Al-Saffar, P. Tuan-Anh, and M. A. A. Almaadeed, “Reinforcement learning-based EV charging and V2G scheduling for renewable energy integration,” IEEE Access, vol. 11, pp. 42810–42826, 2023.

[14] Y. Wang, Q. Chen, T. Hong, and C. Kang, “Review of smart meter data analytics: Applications, methodologies, and challenges,” IEEE Transactions on Smart Grid, vol. 10, no. 3, pp. 3125–3148, 2019.

[15] National Renewable Energy Laboratory (NREL), “Solar Radiation Database (NSRDB),” [Online]. Available: https://nsrdb.nrel.gov, 2023.