energies

MDPI

Article

Dynamic Carbon-Aware Scheduling for Electric Vehicle Fleets Using VMD-BSLO-CTL Forecasting and Multi-Objective MPC

Hongyu Wang [1], Zhiyu Zhao [2], Kai Cui [1], Zixuan Meng [3,*], Bin Li [3], Wei Zhang [1] and Wenwen Li [1]

[1] Marketing Service Center of State Grid Jibei Electric Power Co., Ltd., Xicheng District, Beijing 100051, China; wanghongyu0215@163.com (H.W.); ck99@bupt.cn (K.C.)
[2] State Grid Jibei Electric Power Research Institute, Xicheng District, Beijing 100045, China; jbzcproject6@sina.com
[3] School of Electrical and Electronic Engineering, North China Electric Power University, Changping District, Beijing 102206, China; direfish@163.com
* Correspondence: mengzixuan8398@foxmail.com

Abstract

Accurate perception of dynamic carbon intensity is a prerequisite for low-carbon demand-side response. However, traditional grid-average carbon factors lack the spatio-temporal granularity required for real-time regulation. To address this, this paper proposes a "Prediction-Optimization" closed-loop framework for electric vehicle (EV) fleets. First, a hybrid forecasting model (VMD-BSLO-CTL) is constructed. By integrating Variational Mode Decomposition (VMD) with a CNN-Transformer-LSTM network optimized by the Blood-Sucking Leech Optimizer (BSLO), the model effectively captures multi-scale features. Validation on the UK National Grid dataset demonstrates its superior robustness against prediction horizon extension compared to state-of-the-art baselines. Second, a multi-objective Model Predictive Control (MPC) strategy is developed to guide EV charging. Applied to a real-world station-level scenario, the strategy navigates the trade-offs between user economy and grid stability. Simulation results show that the proposed framework simultaneously reduces economic costs by 4.17% and carbon emissions by 8.82%, while lowering the peak-valley difference by 6.46% and load variance by 11.34%. Finally, a cloud-edge collaborative deployment scheme indicates the engineering potential of the proposed approach for next-generation low-carbon energy management.

Keywords: dynamic carbon factor; electric vehicle scheduling; variational mode decomposition; blood-sucking leech optimizer; model predictive control

Check for updates

Academic Editor: Saim Memon

Received: 23 November 2025
Revised: 9 January 2026
Accepted: 10 January 2026
Published: 16 January 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

The electrification of transport, particularly the proliferation of Electric Vehicles (EVs), is an important component of global decarbonization strategies [1,2]. However, the uncoordinated charging of large-scale EV fleets poses a significant challenge to grid stability and can inadvertently conflict with emission reduction goals [3,4].

A critical requirement for managing this challenge is the transition from static to dynamic carbon accounting. Currently, regulatory practices predominantly rely on annual or regional average emission factors [5–7]. However, this static approach is increasingly incompatible with the evolving requirements of China's "Dual Carbon" strategy. In 2024, the National Development and Reform Commission (NDRC) and the National Energy Administration (NEA) jointly issued the Implementation Opinions on Strengthening the Integration of New Energy Vehicles with the Power Grid. This landmark policy explicitly

mandates the establishment of a comprehensive technical standard system for vehicle-grid interaction (V2G) and emphasizes the necessity for EV loads to actively align with renewable energy generation curves. Furthermore, the Implementation Scheme for Accelerating the Establishment of a Unified and Standardized Carbon Emission Statistical Accounting System further underscores the necessity of high-frequency and traceable data monitoring to ensure the credibility of carbon reduction assessments. Since renewable generation exhibits high volatility, static factors fail to capture the significant spatio-temporal heterogeneity of grid emissions [8,9]. Therefore, utilizing high-resolution dynamic carbon signals to guide EV flexibility has evolved into an important technical pathway for ensuring the environmental effectiveness of demand-side response [10].

A primary obstacle to implementing such guidance is signal latency and computational overhead. Although physics-based carbon flow tracing provides granular data [11–13], its application in real-time control scenarios is severely constrained by inherent calculation bottlenecks. Specifically, global carbon tracing necessitates solving power flow equations for the entire grid topology. As the system scale increases, the computational burden grows exponentially, making it difficult to generate the instantaneous signals required for rapid regulation [14]. Moreover, the conventional centralized calculation and broadcast architecture inevitably introduces transmission delays when interacting with massive distributed edge devices. Consequently, available carbon data is often limited to lagged historical records, failing to serve as the forward-looking signal required for proactive optimization.

Data-driven forecasting provides a compelling solution to this latency problem. However, existing research has yet to achieve a seamless integration between high-precision prediction and real-time scheduling.

First, in the domain of carbon intensity perception, current studies often struggle to handle the complexity of the data. For instance, Yang et al. [15] employed a singular neural network model to predict node carbon factors. While effective for stable trends, this approach lacks the capacity to capture the multi-scale features of highly volatile grid emissions, leading to insufficient accuracy during ramp events. Consequently, advanced hybrid architectures [16,17] are required to decouple these non-stationary sequences.

Second, concerning the scheduling mechanism, existing approaches often exhibit a mismatch between the dynamic nature of the problem and the static methods employed. In the realm of orderly charging, Sha et al. [18] and Chen et al. [19] proposed optimization strategies based on pre-assumed conditions. However, these approaches typically operate as day-ahead global optimizations, generating fixed schedules for the entire operating period. This static perspective overlooks the fact that EV coordination under uncertain carbon signals is fundamentally an "online over-time" problem, a paradigm formally articulated by Duque et al. [20]. Unlike static instances, this paradigm involves limited foresight regarding future inputs and necessitates continuous decision revision. While Duque et al. [20] explored these challenges in the context of combinatorial optimization (e.g., handling job interruptions), we extend this philosophy to the continuous domain of power regulation. Static planning fails to accommodate this reality; once the actual carbon intensity fluctuates, the pre-calculated schedule becomes suboptimal. To implement this "online over-time" processing within a continuous control framework, Model Predictive Control (MPC) offers a robust solution. By utilizing a rolling-horizon mechanism, MPC optimizes decisions over a finite window and continuously updates control actions based on the latest states. This closed-loop feedback effectively handles the "limited foresight" challenge by dynamically adjusting power allocation, offering superior adaptability compared to traditional day-ahead models [21,22].

To address these challenges, this paper proposes an integrated predictive control framework. Specifically, we construct a hybrid forecasting model by integrating VMD

with a predictive core based on CNN, Transformer, and LSTM components (CTL). To address the complex hyperparameter tuning required by this hybrid structure, we apply the Blood-Sucking Leech Optimizer (BSLO) [23] to adaptively identify suitable parameters for each decomposed signal component. This approach is designed to generate high-fidelity, forward-looking carbon signals, which are then utilized as a direct, time-varying objective within a multi-objective MPC strategy. The proposed framework seeks to synergistically balance economic costs, carbon emissions, and grid stability in a computationally tractable, real-time dispatch.

The remainder of this paper is organized as follows: Section 2 details the comprehensive methodology, including the construction of the hybrid VMD-BSLO-CTL forecasting model and the formulation of the multi-objective MPC scheduling strategy based on cloud-edge collaboration. Section 3 presents the experimental results, verifying the accuracy of the carbon factor prediction and demonstrating the effectiveness of the proposed scheduling strategy in reducing costs and emissions under real-world scenarios. Section 4 discusses the economic feasibility, engineering deployment potential, and user incentive mechanisms of the proposed framework. Finally, Section 5 concludes the paper with a summary of key findings and future research directions.

2. Materials and Methods

2.1. Hybrid Carbon Factor Forecasting Model

This Section details the integrated data-driven methodology developed to address the challenges outlined in the introduction. The framework begins with VMD, which is employed to decompose the original complex and non-stationary time series into a set of simpler, more stable intrinsic mode functions (IMFs). Following this, a hybrid predictive core, the CNN-Transformer-LSTM (CTL), is constructed to leverage the respective advantages of each architecture for forecasting the individual components.

Finally, to address the specific challenge of determining optimal hyperparameters for the complex CTL model as it processes the distinct characteristics of each IMF component, the BSLO is applied for adaptive tuning. The principles and implementation of each component within this framework are detailed in the subsequent sections, beginning with VMD.

2.1.1. Signal Decomposition via VMD

To effectively capture the multi-scale dynamics and high volatility of the carbon factor time series, this study employs a hybrid forecasting methodology. We utilize VMD as an advanced and adaptive signal pre-processing technique, chosen for its robustness to noise and its strong mathematical foundation, which effectively avoids common issues like mode-mixing. The method operates by decomposing the original signal f by solving a constrained variational problem:

$$\min_{\{u_k\},\{\omega_k\}} \left\{ \sum_k \left\| \partial_t \left[\left(\delta(t) + \frac{j}{\pi t} \right) * u_k(t) \right] e^{-j\omega_k t} \right\|_2^2 \right\} \quad \text{s.t.} \sum_k u_k = f. \tag{1}$$

In this formulation, the objective is to minimize the cumulative bandwidth of all modes $\{u_k\}$. This bandwidth is estimated by assessing the L^2-norm of the gradient for each mode's analytic signal (obtained via the Hilbert transform) after it has been shifted to its respective baseband center frequency ω_k. This optimization is performed under the hard constraint (s.t.$\sum_k u_k = f$) that the sum of all modes must precisely reconstruct the original signal f. Therefore, VMD serves to implement this decomposition, yielding a discrete set of simpler, more stable Intrinsic Mode Functions (IMFs). Each resulting IMF

is a band-limited component concentrated around a distinct center frequency, making it inherently easier to model. This decomposition-based approach allows our main predictive core to model the unique properties of each signal component separately. These k IMF components are then fed individually into the CNN-iTransformer-LSTM predictive core to generate k separate forecasts. Finally, the overall high-precision carbon factor forecast is reconstructed by summing the individual forecasts from all components.

2.1.2. The CNN-Transformer-LSTM(CTL) Predictive Core

The hybrid predictive core, CNN-Transformer-LSTM (CTL), is illustrated in Figure 1. The model is specifically designed to achieve full-scale feature capture—addressing local, global, and temporal characteristics—from the decomposed IMF components. The data flow and component functions are detailed as follows:

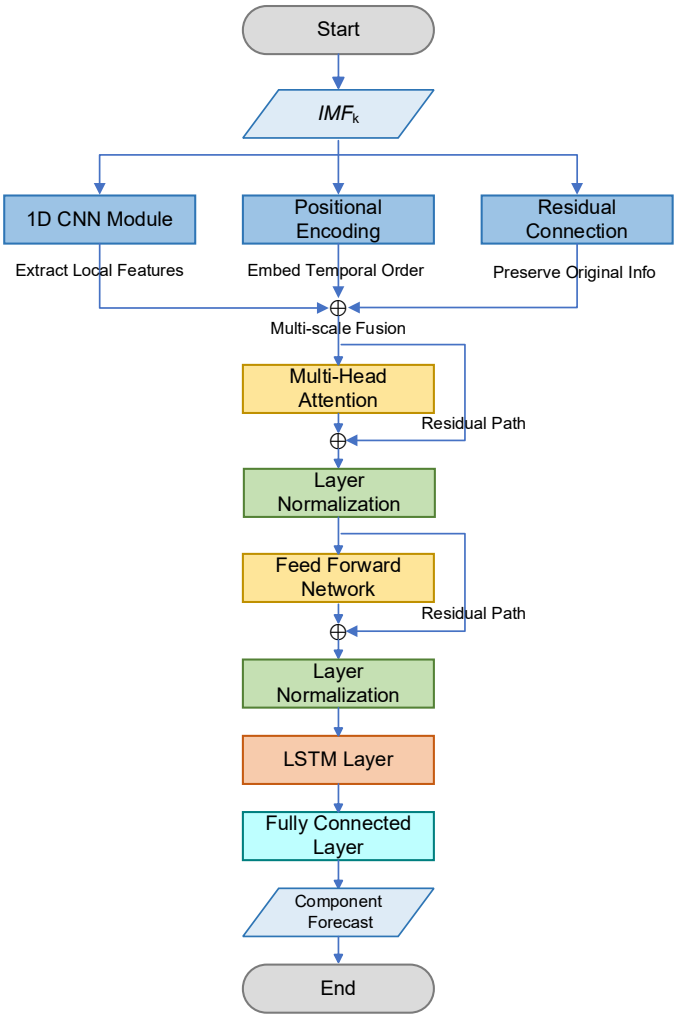


Figure 1. The algorithmic flowchart of the proposed CTL forecasting model.

The input signal (i.e., the IMF_k component from the VMD) first enters a three-branch parallel feature fusion module:

- Local Feature Branch: The signal initially passes through a one-dimensional Convolutional Neural Network (CNN) module. Leveraging the characteristics of convolutional kernels and weight sharing, the CNN block efficiently captures local fluctuations and short-term morphological features within the sequence.

- Positional Information Branch: Concurrently, the original signal is fed into a Positional Encoding unit, which provides the model with essential prior knowledge of the "temporal order" of the sequence.
- Original Information Branch: The signal is also preserved through a Skip-Connection (or residual path).

The outputs from these three branches (local features, positional information, and original information) are then deeply fused via element-wise addition, forming a single, rich feature representation that incorporates multi-scale information.

Subsequently, this fused feature representation is fed into the Transformer Encoder module. The core of the Transformer is the Multi-Head Self-Attention mechanism, which concurrently computes the interdependencies between all time steps in the sequence. This addresses the limitations of a CNN, which can only focus locally, allowing the model to capture long-range dependencies and global contextual patterns within the signal.

Finally, the feature sequence output by the Transformer encoder, which now embeds global context, is passed to a Long Short-Term Memory (LSTM) network. The LSTM utilizes its sophisticated gate mechanisms (forget, input, and output gates) to perform a final temporal dynamic refinement. It smooths and models the dynamic evolution of the global features over time, ultimately passing its hidden state to a fully connected layer to produce the high-precision forecast.

2.1.3. Framework Hyperparameter Optimization: BSLO

As described in Section 2.1.2, the hybrid CTL model possesses a complex architecture with numerous hyperparameters, including the positional embedding dimension, the number of attention heads, the dimensionality of attention keys, the number of LSTM hidden neurons, the dropout rate, and the learning rate. The performance of the entire framework is highly sensitive to these parameters. Manual tuning is subjective and time-consuming, while conventional algorithms like PSO and GWO can easily fall into local optima.

To address this challenge, this study employs the BSLO, a metaheuristic algorithm proposed by Bai et al. [23] in 2024. BSLO is inspired by the foraging and hunting behaviors of blood-sucking leeches in rice fields. The algorithm is modeled on five key strategies: the exploration, exploitation, and switching mechanisms of "directional leeches," the search strategy of "directionless leeches," and a "re-tracking" strategy.

A key feature of BSLO is its dynamic balance between exploration (global search) and exploitation (local search). The algorithm adaptively divides the population (N) into two groups: directional leeches (N_1) and directionless leeches (N_2), governed by Equations (2) and (3):

$$N_1 = \text{floor}\left(N \times \left(m + (1-m) \times \left(\frac{t}{T}\right)^2\right)\right); \quad (2)$$

$$N_2 = N - N_1, \quad (3)$$

where t is the current iteration, T is the maximum iterations, and m is a ratio parameter (set to 0.8 in the original study) that gradually increases the number of directional leeches as the search progresses.

Its core mechanisms rely on these groups:

- Directional Leeches (N_1): These agents simulate the primary hunting force. They dynamically switch between exploration (when far from the prey) and exploitation (when close to the prey), controlled by a "Perceived Distance (PD)" metric.
- Directionless Leeches (N_2): These agents employ a Lévy flight distribution to perform random searches, which enhances population diversity and strengthens the algorithm's ability to escape local optima.

- Re-tracking Strategy: To further prevent stagnation, leeches that are trapped in a local optimum for a set number of iterations are periodically re-initialized, forcing them to search again.

BSLO was selected for this study based on two key criteria. First, on complex benchmark functions (CEC 2017/2019), it has demonstrated competitive performance against conventional algorithms like GWO and PSO. Second, it has been specifically validated for optimizing the parameters of Artificial Neural Networks (ANNs), demonstrating its effectiveness for deep learning applications.

In this study, the BSLO algorithm is applied to find the optimal set of hyperparameters for the CTL model for each IMF component. The objective function is defined as the minimization of the Root Mean Square Error (RMSE) on the validation dataset.

2.2. MPC-Based Multi-Objective EV Optimal Scheduling Framework

Building upon the high-precision carbon factor forecasts developed in the preceding section, this section details the multi-objective optimal scheduling framework designed to utilize these signals. The core challenge is to manage a heterogeneous fleet of EVs in real-time, balancing multiple competing objectives. This chapter first establishes the physical charging and discharging models for the EV fleet. Subsequently, it introduces a MPC strategy as the core framework to synergistically optimize grid stability, user economics, and dynamic carbon emissions.

2.2.1. EV Charging/Discharging Physical Modeling

The EV fleet consists of N heterogeneous vehicles, including a subset $N_{V2G} \subseteq N$ capable of V2G discharging. The physical operation of each vehicle i during its available connection period $[t_i^o, t_i^d]$ is governed by a set of unified power and energy constraints.

First, the charging ($P_{i,t}^{ch}$) and discharging ($P_{i,t}^{dis}$) power at any time t are constrained by their maximum ratings and their mutually exclusive nature. This is formulated using binary variables $u_{i,t}^{ch}$ (charging status) and $u_{i,t}^{dis}$ (discharging status), ensuring a vehicle does not charge and discharge simultaneously:

$$\begin{cases} 0 \le P_{i,t}^{ch} \le u_{i,t}^{ch} \cdot P_i^{char,max} \\ 0 \le P_{i,t}^{dis} \le u_{i,t}^{dis} \cdot P_i^{dis,max} \\ u_{i,t}^{ch} + u_{i,t}^{dis} \le 1 \quad \forall t \in [t_i^o, t_i^d] \\ u_{i,t}^{dis} = 0, \quad \forall i \notin N_{V2G}, \end{cases} \tag{4}$$

where $P_i^{char,max}$ and $P_i^{dis,max}$ are the maximum charging and discharging power ratings for vehicle i. The final constraint ensures that non-V2G vehicles ($i \notin N_{V2G}$) can never discharge.

Second, the battery's State-of-Charge (SOC) dynamics are modeled. The SOC at time t evolves from the previous state $t-1$, and must remain within safe operational limits $[SOC_i^{min}, SOC_i^{max}]$ throughout the dispatch. Furthermore, the vehicle must meet its target energy requirement SOC_i^{target} by its departure time t_i^d:

$$\begin{cases} SOC_{i,t} = SOC_{i,t-1} + \left(\eta^{ch} P_{i,t}^{ch} - P_{i,t}^{dis} / \eta^{dis}\right) \Delta t / E_i^{cap} \\ SOC_i^{min} \le SOC_{i,t} \le SOC_i^{max} \quad \forall t \in [t_i^o, t_i^d] \\ SOC_{i,t_i^d} \ge SOC_i^{target}, \end{cases} \tag{5}$$

where E_i^{cap} is the battery capacity; η^{ch} and η^{dis} are the charging and discharging efficiencies, respectively; and Δt is the uniform time step of the optimization model.

Finally, the total aggregated EV load P_t^{EV} is the sum of all individual vehicle power actions, which, when added to the existing P_t^{Base} (base load), determines the total grid load P_t^{Total}:

$$P_t^{Total} = P_t^{Base} + P_t^{EV} = P_t^{Base} + \sum_{i=1}^{N} (P_{i,t}^{ch} - P_{i,t}^{dis}). \quad (6)$$

2.2.2. MPC-Based Multi-Objective Optimal Scheduling

1. Multi-Objective Formulation and Normalization

The scheduling framework aims to optimize three competing objectives: Grid-Friendliness ($\mathcal{J}_{grid}$), User Economics ($\mathcal{J}_{econ}$), and Environmental Impact ($\mathcal{J}_{carbon}$). To ensure a balanced trade-off and dimensionless comparison, each objective metric ($\mathcal{F}$) is normalized against its respective unoptimized baseline value (B).

- Grid-Friendliness ($\mathcal{J}_{grid}$): This objective co-optimizes two grid metrics: the Load Variance ($\mathcal{F}_{var}$) and the Peak-to-Valley Difference ($\mathcal{F}_{peak}$) of the total power load (P_k^{Total}). This is formulated as a weighted sum of their normalized values over the time horizon $\mathcal{K} = [t, t + H - 1]$:

$$\mathcal{J}_{grid} = w_1 \cdot \frac{\mathcal{F}_{var}}{B_{var}} + w_2 \cdot \frac{\mathcal{F}_{peak}}{B_{peak}}; \quad (7)$$

where the objective components are calculated as:

$$\mathcal{F}_{var} = \sum_{k \in \mathcal{K}} \left(P_k^{Total} - \overline{P}^{Total} \right)^2, \quad \mathcal{F}_{peak} = \max_{k \in \mathcal{K}} \left(P_k^{Total} \right) - \min_{k \in \mathcal{K}} \left(P_k^{Total} \right), \quad (8)$$

Here, P_k^{Total} is the total load at time k, $\overline{P}^{Total}$ is the average total load, and B_{var} and B_{peak} are the baseline values for variance and peak-valley difference, respectively.

- User Economics ($\mathcal{J}_{econ}$): This objective minimizes the total cost for the user ($\mathcal{F}_{cost}$), which is the sum of the energy cost from the grid (based on time-of-use (TOU) price C_k^{price}) and a linearized battery degradation cost (C_i^{deg}) applied only during V2G discharging ($P_{i,k}^{dis}$):

$$\mathcal{J}_{econ} = \frac{\mathcal{F}_{cost}}{B_{cost}}, \quad \mathcal{F}_{cost} = \sum_{k \in \mathcal{K}} \left(C_k^{price} \cdot P_k^{EV} \right) + \sum_{k \in \mathcal{K}} \sum_{i=1}^{N} \left(C_i^{deg} \cdot P_{i,k}^{dis} \right), \quad (9)$$

where P_k^{EV} is the net aggregated power of the EV fleet, C_i^{deg} is the linearized degradation cost per kWh, and B_{cost} is the baseline unoptimized economic cost.

- Environmental Impact ($\mathcal{J}_{carbon}$): This objective minimizes the total physical carbon emissions ($\mathcal{F}_{carbon}$). This is calculated by multiplying the aggregated EV net load P_k^{EV} at each interval by the dynamic carbon emission factor C_k^{factor}:

$$\mathcal{J}_{carbon} = \frac{\mathcal{F}_{carbon}}{B_{carbon}}, \quad \mathcal{F}_{carbon} = \sum_{k \in \mathcal{K}} \left(C_k^{factor} \cdot P_k^{EV} \right), \quad (10)$$

where B_{carbon} is the baseline unoptimized carbon emissions.

2. Model Predictive Control Framework and Constraints

A traditional global optimization approach, which solves the problem for the entire scheduling horizon (T) at once, is computationally intensive and lacks the agility to adapt to real-time changes. Therefore, this study implements a MPC framework. MPC is a

rolling-horizon strategy that makes decisions by optimizing over a finite prediction horizon H into the future, but only implements the decision for the first time step.

At each time step t, the MPC controller solves the following optimization problem over the prediction horizon $k \in [t, t+H-1]$, using the normalized objective components defined above:

$$\min_{P_{i,k}^{char}, P_{i,k}^{dis}} J_{MPC}(t) = \sum_{k=t}^{t+H-1} \left(w_1 \mathcal{J}_{var,k} + w_2 \mathcal{J}_{peak,k} + w_{econ} \mathcal{J}_{econ,k} + w_{carbon} \mathcal{J}_{carbon,k} \right) + J_{terminal}. \quad (11)$$

The optimization is subject to the following constraints:

Standard EV Constraints: Includes EV power limits, SOC boundaries (S_{min} and S_{max}), and the SOC evolution equation for all $k \in [t, t+H-1]$.

Safety-Constrained Dynamic Minimum SOC (Intelligent Hard Constraint):To ensure the long-term feasibility of meeting the departure SOC target (SOC_i^{target}), a dynamic hard constraint is imposed. This constraint proactively calculates the minimum required SOC ($\mathrm{SOC}_{i,k}^{min}$) at every time step k to still reach SOC_i^{target} by the departure time t_i^d, assuming maximum possible charging for the remaining connected duration. This prevents the controller from entering an unrecoverable state:

$$\mathrm{SOC}_{i,k} \geq \max \left(S_{min}, \mathrm{SOC}_i^{target} - \frac{1}{E_i^{cap}} \sum_{j=k}^{t_i^d} P_i^{char,max} \cdot \eta_{char} \cdot W_{i,j} \cdot \Delta t \right), \quad \forall i, \forall k \in \mathcal{K}. \quad (12)$$

Soft Terminal Penalty: The terminal term $J_{terminal}$ is implemented as a soft penalty to guide the system's final state ($t+H$) towards the SOC_i^{target}. The penalty coefficient M is set to 0.5:

$$J_{terminal} = 0.5 \cdot \sum_{i=1}^{N} \max(0, \mathrm{SOC}_i^{target} - \mathrm{SOC}_{i,t+H}). \quad (13)$$

By iteratively resolving this problem at each step t with updated system states, the MPC framework provides an adaptive and computationally tractable solution while guaranteeing long-term safety via the dynamic minimum SOC constraint.

2.3. Experimental Parameter Settings

2.3.1. Computational Environment and Parameter Settings

To comprehensively validate the proposed "Prediction-Optimization" closed-loop framework, we established a high-performance simulation platform. The entire system is developed in MATLAB R2024b. The hybrid forecasting model (VMD-BSLO-CTL) is deployed on a deep learning environment accelerated by an NVIDIA GeForce RTX 4060 Ti GPU, ensuring millisecond-level inference for real-time applications. The subsequent scheduling optimization is modeled using the YALMIP Toolbox (Version 20230622) and solved by the IBM ILOG CPLEX Optimizer (Version 12.10, IBM, Armonk, NY, USA) on an Intel Core i7-12700F CPU (Intel, Santa Clara, CA, USA).

The system is equipped with 32 GB of RAM to efficiently handle high-frequency data processing. The scheduling optimization is formulated as a Quadratic Programming (QP) model. To guarantee global optimality and high-precision convergence, the Optimality Tolerance of the solver is strictly set to 1×10^{-6}. Under this configuration, the prediction inference takes less than 0.05 s, and the average computation time for a single MPC rolling step is approximately 3 s. Considering the 30 min dispatch interval, this computational overhead is negligible, confirming the framework's capability for online real-time operation. The detailed implementation parameters and computational performance metrics are summarized in Table 1.

Table 1. Implementation details and computational environment.

Category	Item	Specification/Value
Hardware	GPU	NVIDIA GeForce RTX 4060 Ti
	CPU	Intel Core i7-12700F
	RAM	32 GB
Software	Environment	MATLAB R2024b
	Modeling Interface	YALMIP Toolbox
	Optimization Solver	IBM ILOG CPLEX Optimizer
Solver Config	Problem Type	Quadratic Programming
	Convergence Criteria	Optimality Tolerance = 1×10^{-6}
Efficiency	Prediction Inference Time	<0.05 s
	MPC Solving Time	Avg. 3.0 s/step

The experimental design is systematically structured into two phases: the validation of the carbon factor prediction model and the application of the multi-objective scheduling strategy.

2.3.2. Forecasting Dataset and Comparative Setup

To verify the robustness and generalization capability of the proposed VMD-BSLO-CTL model, this study utilizes the authoritative Regional Carbon Factor Dataset provided by the UK National Energy System Operator (NESO). The primary dataset covers the period from 1 June 2024, to 31 May 2025, for the South West England region. This area is characterized by a moderate penetration of renewable energy, yielding a carbon factor sequence with appropriate complexity and volatility, which serves as an ideal benchmark for verifying time-series forecasting performance. For data preprocessing, the raw data sampled at 30 min intervals is segmented into look-back windows of 48 time steps, corresponding to a 24 h historical observation period. The dataset is partitioned into training, validation, and testing sets in an 8:1:1 ratio.

Regarding model configuration, the VMD is set to $K = 2$ to separate the series into three distinct sub-signals. Furthermore, the BSLO is employed to eliminate manual tuning bias by adaptively optimizing the key hyperparameters of the predictive network, including the number of attention heads, key dimension, LSTM hidden neurons, and learning rate. To rigorously validate the superiority of the optimization strategy, comparative experiments are conducted against Particle Swarm Optimization (PSO), Bayesian Optimization, and Random Search under identical constraints.

To comprehensively evaluate the prediction performance, the comparative framework is expanded to include two distinct categories: (1) Naïve Baselines (Seasonal Naïve and Persistence models), utilized to verify the model's ability to break through the inertial barrier of the time series; and (2) Advanced Forecasting Models (LSTM, DLinear, and TCN-LSTM-SVM), selected to benchmark deep learning performance. The specific descriptions of these models are detailed in Table 2.

Beyond the primary benchmark, this study further incorporates a supplementary generalization experiment to rigorously test the model's resilience in high-volatility environments. A secondary dataset from the North Scotland region (1 June 2023–31 May 2024) is introduced, which features significantly higher renewable penetration and load fluctuation compared to South West England. Within this extended framework, a multi-step forecasting protocol (covering 1-step, 4-step, and 8-step horizons) is implemented to assess the long-term stability of the VMD-BSLO-CTL model against the comparative algorithms under conditions of increasing uncertainty.

Table 2. Experimental setting of different algorithm group.

Name	Description
Seasonal Naïve	A statistical baseline assuming strict 24 h periodicity, utilizing the value from the previous cycle directly as the forecast.
Persistence	A fundamental benchmark relying on signal inertia, using the current observation as the prediction to test short-term predictability.
LSTM	A classic Recurrent Neural Network (RNN) variant capable of learning long-term dependencies.
DLinear	A state-of-the-art linear model that decomposes time series into trend and seasonal components, known for its efficiency and accuracy.
TCN-LSTM-SVM	An advanced hybrid model combining Temporal Convolutional Networks, LSTM, and Support Vector Machines for robust time-series forecasting.
VMD-BSLO-CTL	The proposed hybrid framework integrating VMD signal decomposition and a BSLO-optimized CTL network for adaptive multi-scale forecasting.

2.3.3. Scheduling Scenario and Fleet Configuration

Following the validation of the prediction model, the high-precision forecast signals are applied to a scheduling scenario based on a region in North China with high renewable energy penetration. The simulation environment is constructed by aggregating real load data from multiple charging stations. Assuming a 20% future schedulable load ratio, the specific composition of the EV fleet was derived by analyzing the peak power and total charging energy characteristics of the real-world load curve. This process reverse-simulated a heterogeneous, station-level EV fleet designed to capture the stochastic nature of user behavior. Specifically, the fleet is stratified into three functional categories: Core Regulation Resources, comprising high-capacity vehicles (50–80 kWh) with fast charging (60 kW) and high-power discharging (15 kW) capabilities; Supplementary Regulation Resources, consisting of long-range vehicles (80 kWh) with standard interfaces for fine-grained regulation; and Rigid and Destination Loads, which include unidirectional fast-charging and slow-charging vehicles that impose mandatory load constraints without participating in discharge.

The scheduling operates under a multi-objective MPC framework employing a rolling horizon mechanism. The weight distribution within the objective function is critically adjusted to reflect real-world trade-offs. In the studied region, high renewable energy penetration combined with Time-of-Use pricing encourages residents to concentrate on charging during midday low-price periods. In this context, aggressively shifting loads to prioritize grid friendliness would likely severely compromise user economic benefits. To navigate this conflict, this study adopts a baseline weight configuration that prioritizes user incentives while maintaining grid stability: Economic Cost is assigned a dominant weight of 0.5, and Carbon Emissions are weighted at 0.2. Meanwhile, Grid Load Variance and Peak-Valley Difference are each assigned a weight of 0.15 to ensure acceptable operational safety.

To comprehensively evaluate the strategy's performance boundaries beyond this baseline, the experimental design is extended to include rigorous verification across two dimensions.

First, standard robustness and sensitivity analyses are conducted to address signal and parameter uncertainties. We implement a robustness test by injecting Gaussian noise into the predicted carbon intensity signals to simulate real-world forecasting errors, ensuring the strategy remains stable under imperfect information. Simultaneously, a sensitivity

analysis is performed by varying objective weights to explore the Pareto frontier between economic and grid indicators.

Second, to specifically address the behavioral uncertainty associated with job interruption (e.g., unexpected early departure), we introduce an Explicit Risk-Averse Mechanism. Unlike the baseline deterministic setting, we incorporate an Individualized Gaussian Hazard Model into the objective function. This model defines a dynamic risk penalty term J_{risk} related to the vehicle's departure proximity:

$$J_{risk} = \sum_{t=1}^{H} \sum_{i=1}^{N_{ev}} \beta_i(t) \cdot \max(0, E_{target,i} - E_{i,t}), \tag{14}$$

where the time-varying risk weight $\beta_i(t)$ follows a Gaussian profile centered at the scheduled departure time $T_{dep,i}$:

$$\beta_i(t) = \lambda_{risk} \cdot \exp\left(-\frac{(t - T_{dep,i})^2}{2\sigma^2}\right) \tag{15}$$

Here, λ_{risk} represents the risk aversion intensity, and σ controls the width of the anxiety window (set to $\sigma = 2$ h). This formulation compels the optimizer to build an energy buffer ("front-loading") as the deadline approaches, minimizing the unserved energy risk. The detailed simulation parameters are listed in Table 3.

Table 3. Simulation parameters for EV scheduling.

Symbol	Description	Value/Unit
Δt	Optimization time step	1 h
SOC_i^{target}	Target SOC at departure	$\geq 0.75 \times E_i^{cap}$
SOC_i^{min}	Minimum State-of-Charge limit	0.1
SOC_i^{max}	Maximum State-of-Charge limit	0.9
η^{ch}/η^{dis}	Charging/Discharging efficiency	0.95
C_i^{deg}	Linearized battery degradation cost	0.0804 CNY/kWh
R^{change}	Battery replacement cost	41,000 CNY

3. Results

3.1. Carbon Factor Forecasting Performance Analysis

1. Signal Decomposition and Adaptive Optimization Analysis

The varying volatility of renewable energy integration introduces significant non-stationarity to the carbon factor sequence. To decouple these complexities, the proposed VMD algorithm decomposes the original time series into modes with distinct central frequencies, as illustrated in Figure 2.

The decomposition results reveal clear physical distinctions among the components. The Fluctuation component (IMF_1), located in the top panel, displays high-frequency oscillations with substantial amplitude, reflecting the rapid intraday dynamics caused by the intermittency of wind and solar power generation. In contrast, the Trend component (IMF_2), shown in the middle panel, exhibits a smooth, low-frequency waveform that captures long-term baseline trends driven by regular generation scheduling and seasonal variations. Finally, the Residual component (res) in the bottom panel isolates high-frequency stochastic noise and random perturbations. This "Divide and Conquer" strategy effectively separates deterministic patterns from stochastic noise.

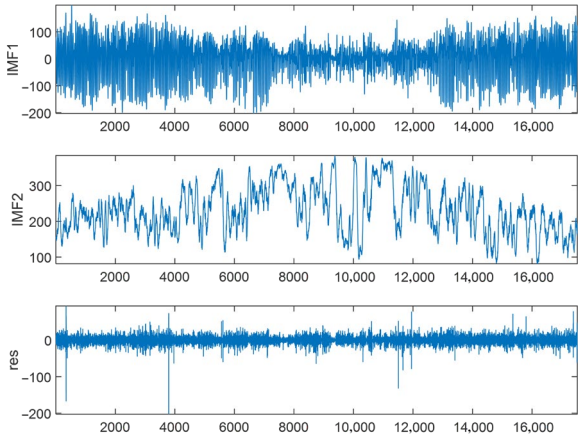


Figure 2. VMD time-domain decomposition diagram.

Based on this decomposition, the BSLO algorithm adaptively tailored the CTL model architecture for each specific component to achieve optimal feature extraction. As detailed in Table 4, for the highly volatile Fluctuation component, BSLO assigned a massive memory capacity (256 LSTM units) and a higher learning rate (4.2×10^{-3}) to track rapid state changes. Conversely, for the smooth Trend component, the optimizer converged to a configuration with a high number of attention heads (32) but minimal LSTM units (8), suggesting a focus on capturing global dependencies via the attention mechanism. Crucially, for the noise-dominated Residual component, the algorithm adopts a minimalist structure (1 head, 4 LSTM units), applying the "Occam's Razor" principle to prevent overfitting stochastic noise. This adaptive configuration significantly contributes to the overall prediction accuracy.

Table 4. Optimal hyperparameters identified by BSLO for each component.

Component	Corresponding Signal	Attention Heads	Key Dimension	LSTM Hidden Units	Learning Rate
Fluctuation	IMF_1	16	32	256	4.2×10^{-3}
Trend	IMF_2	32	64	8	1.3×10^{-3}
Residual	*res*	1	2	4	1.0×10^{-3}

2. Comparative Justification of Optimization Strategy

To evaluate the effectiveness of the proposed BSLO algorithm in hyperparameter optimization, we conducted a comparative experiment against three standard methods: Particle Swarm Optimization (PSO), Bayesian Optimization (TPE), and Random Search.

A subset of the operating data from 1 May 2025, to 31 May 2025 (approximately 1500 samples) was selected for this analysis. The data was decomposed into three distinct sub-signalsvia VMD: IMF 1 (Daily Periodic Component), IMF 2 (Trend Component), and Residual (High-Frequency Component). To ensure a consistent comparison under limited computational resources, the population-based algorithms (BSLO and PSO) were configured with a population size of $N = 10$ and a maximum of $T = 50$ iterations. Similarly, Bayesian Optimization was limited to 50 function evaluations. The objective was to minimize the Mean Squared Error (MSE) on the validation set.

Figure 3 presents the convergence curves of the different algorithms.

- Daily Periodic Component (IMF 1): IMF 1 represents the dominant 24 h cyclic pattern of the carbon factor. As shown in Figure 3a, BSLO (solid red line) demonstrates efficient search capability, achieving a final MSE of 0.0380, which is lower than both

the Random Search baseline (0.0415) and PSO (0.0386). This indicates BSLO's superior ability in capturing the primary periodicity of the data.

- Trend Component (IMF 2): IMF 2 reflects the long-term evolutionary trend. In Figure 3b, the PSO algorithm (dashed blue line) tends to flatten around iteration 15 (MSE $\approx$ 0.0395), suggesting a potential local optimum. In contrast, BSLO continues to optimize throughout the process, reaching a lower error of 0.0350, proving its robustness in tracking evolutionary trends.
- High-Frequency Component (Residual): The Residual contains stochastic fluctuations and noise. For this complex component shown in Figure 3c, BSLO achieved the best performance with a final MSE of 0.7028, significantly outperforming Bayesian Optimization (0.7811) and Random Search (0.7593).

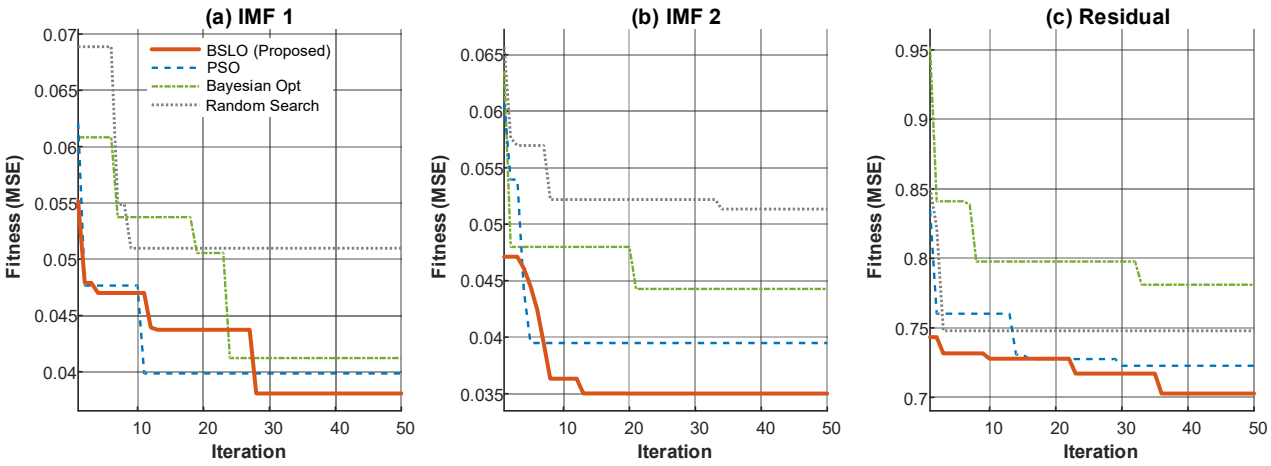


Figure 3. Comparative convergence curves of the BSLO algorithm versus PSO, Bayesian Optimization, and Random Search across three VMD-decomposed components: (**a**) Daily Periodic Component (IMF 1), (**b**) Trend Component (IMF 2), and (**c**) High-Frequency Component (Residual).

The results indicate that under the same number of iterations, BSLO exhibits stable convergence and superior optimization ability across periodic, trend, and stochastic components compared to the benchmarked methods. This justifies the utilization of BSLO for the model's hyperparameter tuning.

3. Ablation Study of Key Components

To justify the complexity of the proposed pipeline and quantify the contribution of each key component, we conducted an ablation study. We designed three variants of the model for comparison:

- CTL: The standalone deep learning model (CNN-Transformer-LSTM) without decomposition or optimization, serving as the base predictor.
- CTL-BSLO: The base predictor optimized by BSLO, but without signal decomposition (raw data input).
- VMD-CTL: The predictor with VMD, but using fixed hyperparameters (without BSLO optimization).

The results across 1-step, 4-step, and 8-step horizons are presented in Table 5.

As shown in Table 5, both VMD and BSLO contribute significantly to the prediction accuracy:

- Impact of Decomposition (VMD): Comparing the CTL and VMD-CTL variants, the introduction of VMD results in a drastic reduction in error. For the 1-step horizon, the MAPE decreases from 21.59% (CTL) to 12.16% (VMD-CTL). This confirms that

decomposing the non-stationary carbon intensity signal into stable sub-modes is the primary factor in improving model performance.

- Impact of Optimization (BSLO): Comparing VMD-CTL with the Proposed (VMD-BSLO-CTL) model (MAPE 9.15%), the BSLO optimization further reduces the MAPE by approximately 3.01%. This validates that manual or fixed parameter settings are insufficient for the diverse sub-signals generated by VMD, and adaptive optimization is necessary to fully unlock the model's potential.
- Synergistic Effect: The proposed framework, which combines both components, achieves the lowest RMSE and MAPE across all horizons. This demonstrates that the complexity introduced by the VMD-BSLO pipeline translates directly into tangible performance gains.

Table 5. Ablation study results of different model variants.

Horizon	Metric	CTL	CTL-BSLO	VMD-CTL	Proposed
1 Step	MAPE (%)	21.59%	14.94%	12.17%	9.15%
	RMSE	19.67	17.71	11.88	10.15
4 Steps	MAPE (%)	33.97%	27.74%	19.26%	14.80%
	RMSE	30.69	29.34	19.09	17.10
8 Steps	MAPE (%)	46.10%	35.66%	22.51%	18.39%
	RMSE	39.74	38.38	25.25	23.35

4. Comprehensive Evaluation of Prediction Performance

The final prediction results on the test set were obtained by aggregating the forecasts of the three decomposed components.

Figure 4 illustrates the single-step prediction performance. Visually, the proposed model exhibits a strong capability to track the ground truth curve, accurately capturing both the smooth transitions during valley periods and the sharp inflection points during ramping events. While minor deviations persist at extreme peaks due to the inherent stochasticity of high-frequency noise, the overall fitting degree remains exceptionally high.

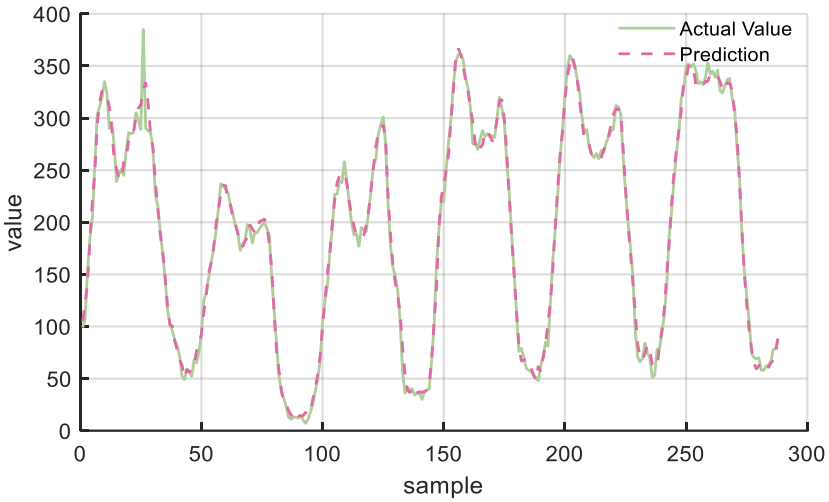


Figure 4. Comparison of single-step prediction and ground truth.

To further evaluate the consistency and error distribution of the model, Figure 5 presents the density scatter plot of the predicted values against the ground truth. The color gradient indicates the data point density, where dark purple represents high-density regions and dark orange denotes low-density areas. It is evident that the vast majority of data points

are tightly clustered along the ideal $y = x$ diagonal, yielding a coefficient of determination R^2 of 0.991. This high goodness of fit confirms that the model has captured the underlying physical patterns without significant systematic bias. Although a few scattered outliers exist in the low-density orange regions, the overall distribution demonstrates the robustness of the model across different carbon intensity levels. Quantitatively, the model achieves an RMSE of 10.15 and a Mean Absolute Percentage Error of 9.15% for the single-step horizon.

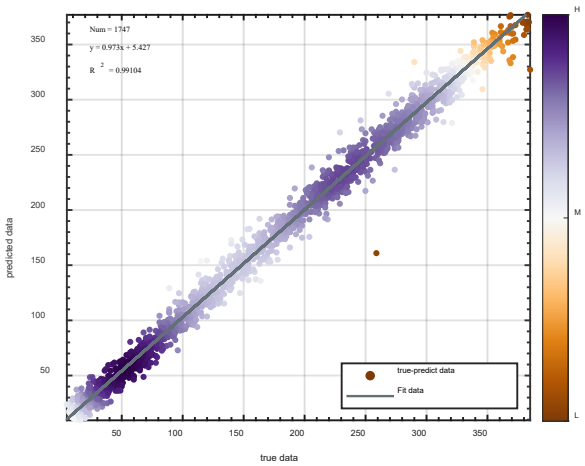


Figure 5. Density scatter plot of prediction results.

To explicitly quantify the model's prediction reliability, we conducted probabilistic forecasting using the Bootstrap residual method. As shown in Figure 6, the 95% confidence interval (blue shaded area) tightly envelopes the ground truth curve. The narrow bandwidth of the confidence interval indicates that the model maintains high certainty and low variance even during peak fluctuations, further validating the stability of the BSLO-optimized parameters.

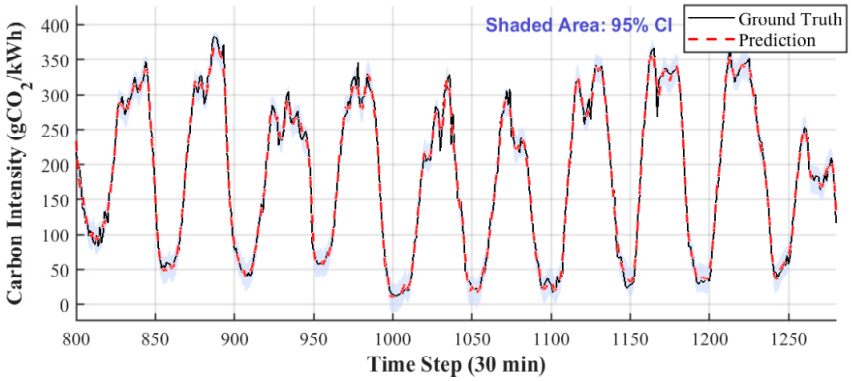


Figure 6. Probabilistic forecasting results with 95% confidence intervals based on the Bootstrap method.

To validate the model's reliability for the MPC rolling horizon, the prediction performance was evaluated across three different time steps: 1 step (30 min), 4 steps (2 h), and 8 steps (4 h). Table 6 benchmarks the proposed VMD-BSLO-CTL model against two distinct categories: (1) Naïve Baselines, including Seasonal Naïve and Persistence models, introduced to rigorously test against the data's inherent predictability; and (2) Advanced Forecasting Models, including LSTM, TCN-LSTM-SVM, and DLinear.

Table 6. Performance comparison across different prediction horizons.

Horizon	Metric	Seasonal Naive	Persistence	LSTM	DLinear	TCN-LSTM-SVM	Proposed
1 Step	MAPE (%)	53.31%	12.71%	28.77%	11.47%	10.09%	9.15%
	RMSE	61.57	21.70	21.44	18.77	17.11	10.15
4 Steps	MAPE (%)	53.35%	44.88%	39.29%	31.18%	22.13%	14.80%
	RMSE	61.61	64.04	33.28	40.78	28.36	17.1
8 Steps	MAPE (%)	53.40%	102.14%	51.08%	55.90%	32.28%	18.39%
	RMSE	61.67	108.87	41.88	53.37	37.39	23.35

The quantitative results indicate the performance stability of the proposed framework, particularly in multi-step forecasting:

- Short-term Inertia vs. True Learning: At the 1-step horizon, the Persistence model achieves a remarkably low MAPE of 12.71%. This confirms the high inertial autocorrelation of the grid carbon intensity. However, the proposed model effectively breaks through this "inertia barrier," further reducing the MAPE to 9.15%, demonstrating its ability to capture high-frequency fluctuations that simple autoregression misses.
- Long-term Robustness: The contrast becomes stark as the horizon extends. The Seasonal Naïve model shows a consistently high error (>53%), indicating the lack of simple 24 h periodicity. More critically, the performance of the Persistence model collapses at the 8-step horizon, with MAPE skyrocketing to 102.14% and RMSE to 108.87. In comparison, the proposed model demonstrates exceptional robustness, maintaining a MAPE of 18.39% even at the 4 h horizon. While other deep learning models (e.g., TCN-LSTM-SVM at 32.28%) also degrade, the proposed decomposition-ensemble strategy effectively mitigates error accumulation, ensuring reliable forward-looking signals for scheduling optimization.

To further substantiate the generalization capability of the proposed framework, we conducted an additional evaluation using the North Scotland dataset, a region characterized by high wind penetration and extreme load volatility. Table 7 presents the performance comparison across varying prediction horizons (1-step, 4-step, and 8-step), utilizing RMSE, MAE, and R^2 as key metrics.

Table 7. Generalization performance comparison on North Scotland region.

Horizon	Metric	Seasonal Naive	Persistence	LSTM	DLinear	TCN-LSTM-SVM	Proposed
1 Step	MAE	16.44	3.43	5.43	2.57	4.58	3.75
	RMSE	39.43	14.03	14.93	11.05	15.17	10.95
	R^2	−0.676	0.788	0.761	0.832	0.753	0.871
4 Steps	MAE	16.43	7.23	7.51	5.63	6.59	4.74
	RMSE	39.45	25.67	20.67	21.26	21.2	14.53
	R^2	−0.675	0.291	0.593	0.377	0.571	0.86
8 Steps	MAE	16.43	9.55	8.96	9.25	8.46	5.74
	RMSE	39.49	29.96	23.87	26.21	24.95	16.61
	R^2	−0.675	0.036	0.456	0.054	0.383	0.828

The results reveal a critical divergence in model performance as the prediction horizon extends:

Short-term Competitiveness: In the single-step (1-Step) scenario, linear-based models such as DLinear and the Persistence baseline demonstrate competitive performance, particularly in MAE (2.57 gCO_2/kWh and 3.43 gCO_2/kWh, respectively). This is attributed to the strong autocorrelation inherent in the high-resolution sampling, which simple autoregressive mechanisms can exploit.

Long-term Degradation of Baselines: However, as the horizon increases to 4 and 8 steps, these benchmark models suffer from catastrophic degradation due to the accumulation of recursive errors. For instance, the R^2 of DLinear plummets from 0.832 (Step 1) to a mere 0.054 at Step 8, indicating an inability to capture long-term dependencies in such a volatile environment. Similarly, standard deep learning models like LSTM and TCN-LSTM-SVM show significant performance decay, with RMSEs rising above 23 gCO_2/kWh.

Superior Robustness of Proposed Model: In sharp contrast, the proposed VMD-BSLO-CTL strategy demonstrates exceptional long-term stability. Even at the 8-step horizon, the model maintains an RMSE of 16.61 gCO_2/kWh and an R^2 of 0.828, significantly outperforming the second-best model. This resilience confirms that the VMD effectively isolates high-frequency noise, while the BSLO-optimized architecture accurately models the complex non-linear features that simpler linear models fail to capture. This verifies that the proposed method is not only accurate for immediate dispatch but also highly reliable for longer-term look-ahead scheduling in complex energy systems.

3.2. MPC-Based Scheduling Optimization Analysis

The performance of the proposed scheduling strategy is evaluated through a comprehensive multi-dimensional assessment, as summarized in Figure 7. This comparative analysis reveals the complex trade-offs inherent in high-renewable and time-of-use environments. As shown by the blue dotted line in Figure 7, the Single-Objective MPC strategy, which strictly prioritizes grid stability (minimizing load variance and peak-valley difference), yields the smoothest load curve. However, this comes at a severe penalty: to flatten the grid curve, the algorithm forces EV charging into expensive or high-carbon time windows that do not align with the optimal midday solar generation. Consequently, its economic cost and carbon emissions actually exceed those of the uncoordinated baseline. In contrast, the proposed Multi-Objective MPC achieves an optimal balance. By dynamically weighting user incentives against grid needs, it successfully captures the "green and cheap" charging windows. Compared to the unoptimized baseline, the proposed framework reduces economic costs by 4.17% and carbon emissions by 8.82%, while simultaneously reducing the peak-valley difference by 6.46% and load variance by 11.34%, demonstrating that appropriate scheduling can unlock economic value while supporting the grid.

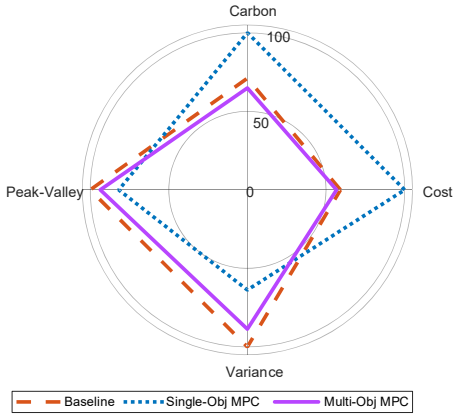


Figure 7. Multi-dimensional performance comparison.

To analyze the performance enhancement mechanism in depth under the multi-objective strategy, Figure 8 illustrates the macroscopic effect of load component redistribution over time, while Figure 9 reveals the micro-dispatch mechanism behind the fleet's collective benefits. Unlike the baseline scenario, where charging load merely overlaps with existing peaks, the dispatchable load under the multi-objective strategy demonstrates highly intelligent time-shifting behavior and dynamic objective prioritization:

- Grid Support during Morning Peak: During the 06:00–07:00 base load morning peak, the optimizer prioritizes the essential peak-shaving objective over instantaneous cost. This compels V2G vehicles to execute net discharge (red blocks in Figure 8). Figure 9 provides the micro-validation of this response: core V2G vehicles (e.g., dark traces) actively initiate SOC decline during this early critical grid period.
- Coordinated Green Absorption and Conflict Navigation (11:00–14:00): Subsequently, the charging load is clustered within this optimal window. This period offers low-carbon intensity and economic favorability, but the concentrated charging also presents a potential high-stress zone for the grid. The coordinated SOC trajectories (Figure 9) confirm that vehicles utilize this cheap, low-carbon midday window for rapid replenishment after their morning discharge, while also actively avoiding the charging peak. The successful dispatch proves that the strategy actively managed this stress, ensuring the cumulative net load curve (dark blue line in Figure 8) remains stable while maximizing green energy absorption.

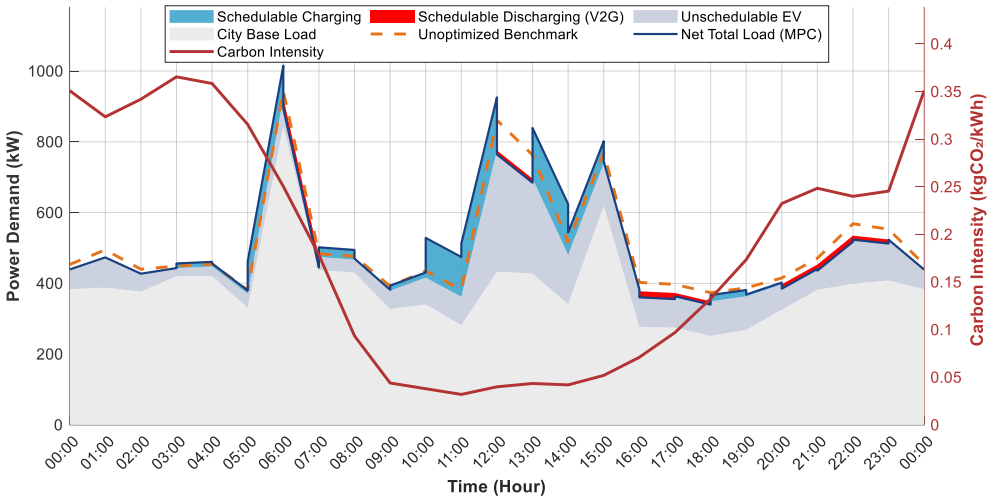


Figure 8. Load profile components and dynamic optimization.

This precise coordination, adhering strictly to safety constraints and final energy requirements, validates the strategy's ability to achieve simultaneous optimization across economic, grid, and environmental objectives.

Furthermore, to address the challenge of uncertainty in real-world operations—where perfect forecasting is unattainable—we evaluated the robustness of the proposed framework against prediction errors. A robustness test was conducted by introducing Gaussian noise with a margin of $\pm 15\%$ to the input carbon factor sequence during the optimization phase. Crucially, while the MPC decisions were made based on these noisy forecasts, the final performance metrics were calculated using the actual (ground truth) carbon intensity and electricity prices. As presented in Table 8, the system demonstrates strong resilience to input uncertainty. Despite the significant noise, the performance degradation is minimal: the deviation in economic cost is only +0.19%, and the increase in carbon emissions is limited to +1.01%. These results indicate that the proposed multi-objective optimization

framework maintains high performance stability and engineering reliability even when operating with imperfect information.

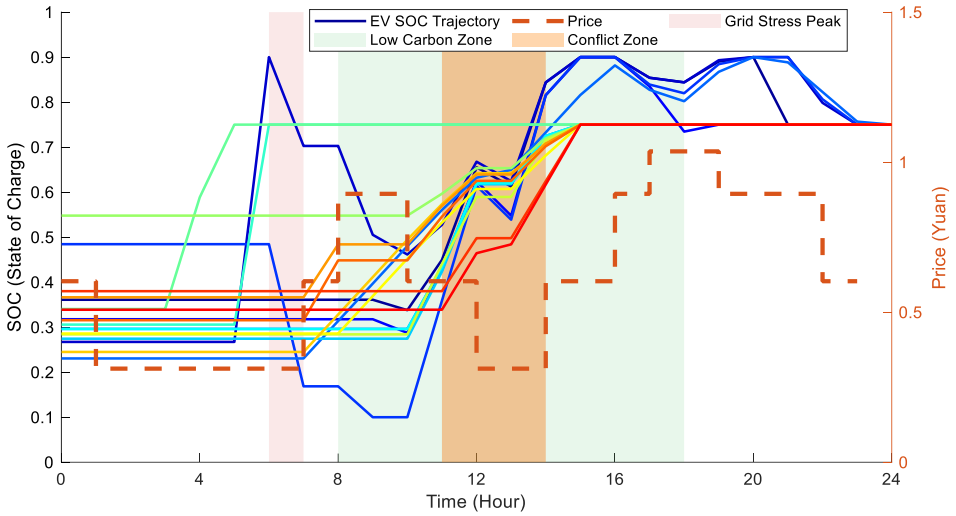


Figure 9. Heterogeneous fleet SOC trajectories and micro-dispatch. Note: The solid lines with various colors represent the individual SOC trajectories of different EVs in the fleet.

Table 8. Robustness Test Results under ±15% Prediction Noise.

Metric	Ideal Case	Robust Case	Deviation (%)
Economic Cost	1591.92 CNY	1594.95 CNY	0.19%
Carbon Emissions	355.65 $kgCO_2$	359.25 $kgCO_2$	1.01%
Peak-Valley Diff	533.99 kW	549.65 kW	2.93%
Load Variance	5.20×10^5	5.07×10^5	−2.33%

Having established the system's resilience to external data uncertainty, it is equally critical to examine the impact of internal parameter selection, specifically the objective weights. To evaluate the robustness of the proposed strategy against weight variations and understand system behavior under varying preferences, a sensitivity analysis was conducted based on the Pareto Frontier. We categorized the optimization objectives into two conflicting groups: "Operational Costs" (Economic Cost and Carbon Emissions) and "Grid Stability" (Peak-Valley Difference and Load Variance). Within these groups, fixed internal ratios were maintained to reflect specific scenario characteristics: a 5:2 ratio was applied between economic cost and carbon emissions to represent their relative market importance, while a 1:1 ratio was assigned to peak-valley difference and load variance as they equally describe grid smoothness.

A preference parameter $\lambda \in [0.1, 0.9]$ was then introduced as a weighting lever to dynamically adjust the trade-off between the two groups. Specifically, the weight of the "Operational Costs" group is scaled by λ, while the "Grid Stability" group is scaled by $(1 - \lambda)$. As λ increases, the optimizer progressively prioritizes cost and emission reductions over grid flatness. Figure 10 illustrates the resulting trends, plotting Economic Cost (a) and Carbon Emissions (b) against Grid Stress.

The results in Figure 10 reveal clear patterns:

- Clear Trade-off: A natural conflict exists between minimizing operational costs and maintaining a perfectly flat grid load. Reducing costs and emissions drives the system to cluster EV charging during low-price, low-carbon midday windows. While operationally efficient, this concentration reduces the degree of load leveling compared to a strict grid-prioritized strategy. Conversely, forcing a perfectly flat load curve requires

shifting demand to expensive or high-carbon periods, leading to a sharp increase in costs.

- Optimal Balance: The strategy selected in this paper (marked by the red star, $\lambda = 0.7$) is located at the "knee point" of the Pareto curve. This point represents the system's "sweet spot," achieving the majority of potential cost and carbon savings while maintaining grid stress at an acceptable level. Deviating from this point would either result in diminishing returns in savings or cause a disproportionate penalty in grid stability, confirming that the chosen parameters capture the most efficient operating point.

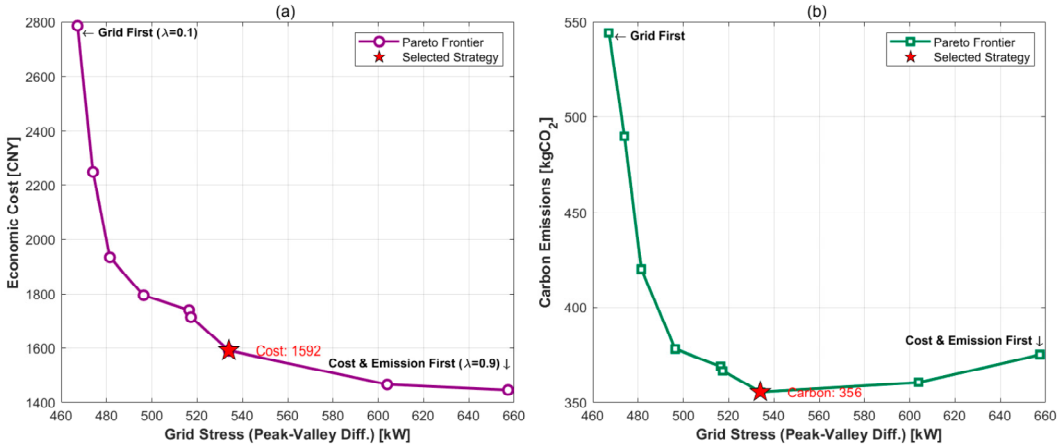


Figure 10. Sensitivity analysis of scheduling preferences. (**a**) Trade-off between economic cost and grid stability; (**b**) trade-off between carbon emissions and grid stability. The red star indicates the selected strategy ($\lambda = 0.7$). Note: The arrows indicate the specific parameter values ($\lambda = 0.1$ and $\lambda = 0.9$) for the endpoints of the Pareto frontier.

Finally, to complete the robustness verification, we extend the evaluation from "data and parameters" to the critical dimension of "user behavior uncertainty," specifically the risk of job interruption (e.g., unexpected early departure). While the previous tests validated resilience against signal noise and weight variations, real-world operations must also account for users disconnecting earlier than scheduled. To address this, an additional comparative test was conducted using the proposed Risk-Averse Robust MPC, which incorporates an explicit Gaussian hazard penalty to model departure anxiety.

Table 9 presents the performance deviation of this robust strategy compared to the baseline. The results reveal a distinct "Front-Loading" mechanism: the system actively builds a "Safety Energy Buffer" by accelerating charging 2–3 h before the scheduled departure. As shown in the table, this active defense capability incurs a marginal "Price of Robustness": the economic cost increases by only 0.81% (1591.92 to 1604.87 CNY), and carbon emissions rise by 1.57%. Notably, this safety enhancement does not compromise grid stability; the peak-valley difference remains unchanged, while load variance actually improves slightly (-2.48%) due to the smoother distribution of charging power across the safety window. This confirms that the proposed framework can secure high service reliability against behavioral uncertainty with negligible economic trade-offs.

Table 9. Performance comparison under behavioral uncertainty (job interruption risk).

Metric	Ideal Case	Robust Case	Deviation (%)
Economic Cost	1591.92 CNY	1604.87 CNY	+0.81%
Carbon Emissions	355.65 $kgCO_2$	361.23 $kgCO_2$	+1.57%
Peak-Valley Diff	533.99 kW	533.99 kW	0.00%
Load Variance	5.20×10^5	5.07×10^5	-2.48%

4. Discussion

The proposed framework is designed with practical engineering viability at its core, contextualized within a major demonstration project by the State Grid Corporation of China. Beyond the theoretical performance verified in the simulation, the practical deployment involves specific considerations regarding system architecture, operational granularity, and economic feasibility.

1. Cloud-Edge Collaborative Architecture and Hardware Implementation

The deployment strategy is specifically designed to address the spatio-temporal precision deficiencies of traditional carbon metering methods. Traditional grid-level carbon accounting, constrained by computationally intensive power flow analysis, is often limited to coarse temporal resolutions and subject to transmission delays. This resolution is insufficient for the fine-grained guidance required for EV charging regulation. To overcome this, our framework utilizes an integrated "cloud-edge" deployment scheme:

- Cloud Platform Layer (Carbon Tracing): The central cloud platform focuses on macroscopic carbon flow tracing. It performs global power flow analysis to calculate the baseline carbon intensity of the regional grid. This baseline data is then disseminated to edge gateways in real-time.
- Edge Computing Layer (Local Calculation and Prediction): The edge gateway, deployed at the charging station, acts as the intelligent local node. It integrates the downlinked grid carbon data with real-time operational data from local distributed resources (such as photovoltaics and energy storage systems). By synthesizing these inputs, the gateway calculates a precise local carbon factor. Subsequently, the lightweight VMD-BSLO-CTL prediction model deployed on the gateway performs high-frequency forecasting of this local carbon intensity. These generated predictive signals can then serve as the critical forward-looking inputs required for the multi-objective scheduling framework, effectively bridging the gap between carbon monitoring and active grid regulation.

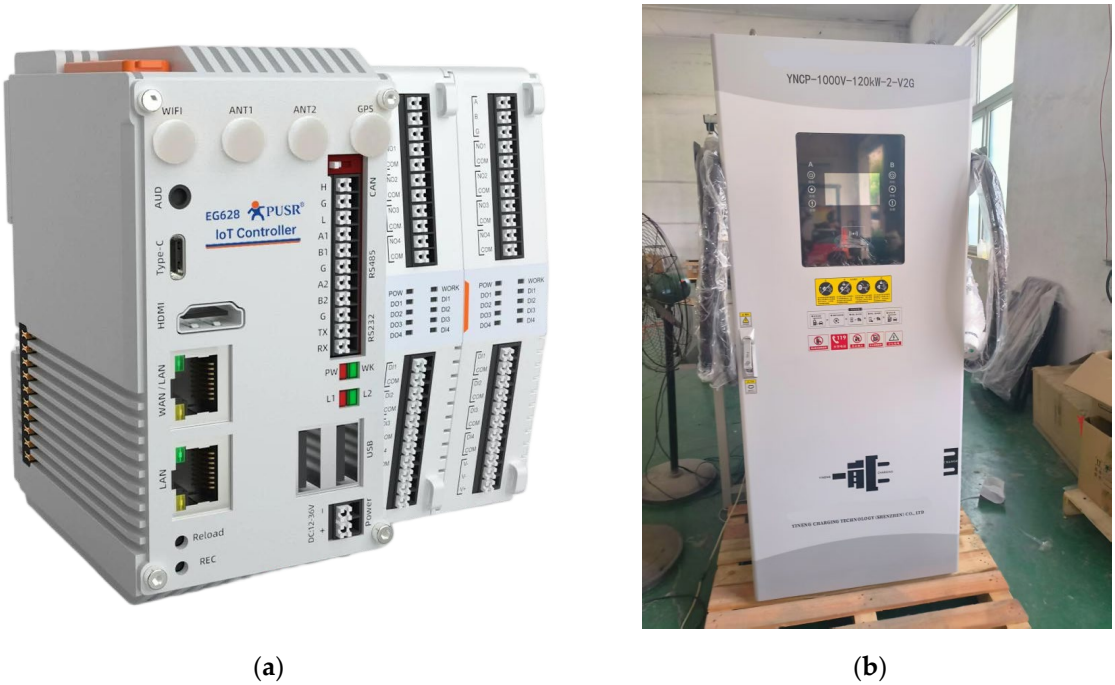
Our pilot hardware selection confirms the physical feasibility of this edge layer. We utilized an industrial-grade edge computing gateway (USR-EG628-G4, sourced from USR IOT, Jinan, China) powered by an ARM Cortex-A9 processor (1.2 GHz). The device is equipped with 16 GB of DDR3 memory and 64 GB of eMMC storage, running a customized Linux OS based on Ubuntu 18.04. In terms of connectivity, the gateway integrates 4G/5G modules, ensuring stable data transmission with the upstream cloud server. The gateway communicates directly with intelligent charging piles via CAN bus or Modbus to collect operational status. This configuration confirms that the prediction model and logic can be executed efficiently on standard industrial hardware.

Figure 11 presents the core hardware equipment selected for the demonstration project. Figure 11a shows the industrial edge gateway (USR-EG628-G4) which hosts the lightweight prediction model, while Figure 11b depicts the 120 kW V2G charging unit that executes the bidirectional dispatch commands.

2. Data Granularity and Operational Synchronization

A key engineering decision in this study was determining the appropriate time granularity for "real-time" control. While modern smart meters are capable of minute-level sampling, the bottleneck lies in the upstream data source: regional grid operators typically calculate and release carbon intensity data at 1 h intervals due to the complexity of the calculation. To ensure our scheduling decisions are based on valid inputs rather than artificial interpolations, we synchronized our control frequency with this upstream data rhythm. Consequently, the system adopts a 1 h rolling horizon. This alignment serves a

critical practical purpose: it prevents the charging equipment from switching too frequently. If we were to adopt a minute-level control scheme as theoretically possible, the frequent start-stop actions would cause excessive mechanical wear on the relays and switches inside the charging piles. Therefore, aligning the scheduling interval with the 1 h data update is a pragmatic choice that balances control accuracy with equipment longevity.



(**a**) (**b**)

Figure 11. Core hardware components for the edge-cloud deployment: (**a**) The intelligent edge computing gateway (USR-EG628-G4) deployed for local carbon factor prediction; (**b**) The 120 kW integrated V2G charging station unit capable of bidirectional power flow.

3. Retrofit Economics and User Alignment

Economically, the framework offers significant advantages by reducing retrofit costs. Rather than requiring a complete overhaul of charging stations, the system functions as an intelligent add-on. By deploying the edge gateway as a "plug-in" module, existing charging piles can be upgraded to possess carbon-aware capabilities. This "smart upgrade" approach offers a cost-effective pathway for station operators compared to replacing entire infrastructure. Operationally, the station-level scheduling is designed to be compatible with user interests. The optimization occurs strictly within the user's connected parking duration, shifting the charging load temporally from high-carbon windows to low-carbon windows. This alignment allows the grid-friendly regulation to naturally translate into economic benefits for users through Time-of-Use arbitrage, creating a sustainable incentive loop.

4. Limitations and Future Directions

While the proposed Multi-Objective MPC framework provides a robust solution, we acknowledge that real-world EV scheduling involves stochastic uncertainties. As noted in the literature regarding combinatorial optimization under limited foresight (e.g., Duque et al. [20]), future iterations of this system could benefit from integrating advanced online algorithms that better account for job interruptions and extreme probability events. Additionally, future work will focus on validating the scalability of the framework in large-scale fleet coordination.

5. Conclusions

This paper proposes a closed-loop framework integrating high-precision carbon factor forecasting with multi-objective EV scheduling to address the challenges of dynamic carbon accounting and grid regulation. The key conclusions are as follows:

- Robust Forecasting Engine: A hybrid VMD-BSLO-CTL model was constructed to predict dynamic carbon factors. By decomposing non-stationary sequences and adaptively optimizing hyperparameters, the model achieved a MAPE of 9.15% on the UK National Grid dataset, demonstrating superior robustness in multi-step forecasting compared to traditional baselines.
- Synergistic Optimization: A multi-objective MPC strategy was developed to navigate the trade-offs between economic incentives and grid stability. Simulation results indicate that the proposed strategy achieves a superior overall balance compared to the uncoordinated baseline. The framework successfully achieved improvements across all four key performance indicators: reducing economic costs by 4.17% and carbon emissions by 8.82%, while simultaneously lowering the peak-valley difference by 6.46% and load variance by 11.34%. These results validate that the multi-objective approach can effectively unlock economic and environmental value while supporting grid smoothness.
- Engineering Potential: The pilot implementation on industrial edge gateways (ARM Cortex-A9, 1.2 GHz) indicates the potential feasibility of the proposed "Cloud-Edge" architecture. The results suggest that high-precision carbon-aware scheduling is likely achievable through a cost-effective retrofit strategy, utilizing existing 4G/5G networks and edge intelligence to bridge the gap between macroscopic grid data and local regulation needs.

Author Contributions: Conceptualization, H.W., Z.Z., K.C. and B.L.; methodology, Z.M.; software, Z.M.; validation, Z.M.; formal analysis, Z.M. and B.L.; investigation, Z.M., W.Z. and W.L.; resources, Z.Z. and K.C.; data curation, Z.M., W.Z. and W.L.; writing—original draft preparation, Z.M.; writing—review and editing, Z.M.; visualization, Z.M.; supervision, H.W.; project administration, H.W.; funding acquisition, H.W. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the State Grid Corporation of China Science and Technology Project, grant number B3018524000Q.

Data Availability Statement: The data presented in this study are available in part. The carbon intensity data analyzed in this study are publicly available and can be found at the NESO Data Portal (https://www.neso.energy/data-portal/regional-carbon-intensity-forecast, accessed on 10 November 2025). The real-world electric vehicle scheduling and load data presented in this study are proprietary and are not publicly available due to privacy and commercial confidentiality restrictions.

Acknowledgments: We are very grateful for the support and cooperation of North China Electric Power University. During the preparation of this manuscript, the authors used ChatGPT-4o (OpenAI, San Francisco, CA, USA) for the purposes of improving language readability and polishing the text. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: Authors Hongyu Wang, Kai Cui, Wei Zhang and Wenwen Li were employed by Marketing Service Center of State Grid Jibei Electric Power Co., Ltd. Author Zhiyu Zhao was employed by State Grid Jibei Electric Power Research Institute. The authors declare that this study received funding from the State Grid Corporation of China. The funder had the following involvement with the study: authors Hongyu Wang, Zhiyu Zhao, Kai Cui, Wei Zhang, and Wenwen Li are employees of the funding agency and participated in the study design, data collection, and decision to publish. The remaining authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

EV	Electric Vehicle
V2G	Vehicle-to-Grid
VMD	Variational Mode Decomposition
BSLO	Blood-Sucking Leech Optimizer
CTL	CNN-Transformer-LSTM
MPC	Model Predictive Control
NESO	National Energy System Operator
SOC	State of Charge
TOU	Time-of-Use
KPI	Key Performance Indicator
MAPE	Mean Absolute Percentage Error
RMSE	Root Mean Square Error

References

1. Yi, J.; Lin, W.; Ren, X.; Cui, Y.; Han, N. Research on construction of index system and comprehensive evaluation method of development level of new power system. *Power Syst. Technol.* **2024**, *48*, 3758–3768. [CrossRef]
2. Luo, J.; Zhuo, W.; Liu, S.; Xu, B. The optimization of carbon emission prediction in low carbon energy economy under big data. *IEEE Access* **2024**, *12*, 14690–14702. [CrossRef]
3. Yu, X.; Ji, Z.; Ji, L.; Niu, D.; Xu, X. Analysis on carbon emission reduction potential of electric vehicles in China under goal of carbon neutrality and carbon peaking. *Smart Electr. Power* **2024**, *52*, 25–31.
4. Li, J.; Guo, Z.; Ma, S.; Zhang, G.; Wang, H. Overview of the "source-grid-load-storage" architecture and evaluation system under the New Power System. *High Volt. Eng.* **2022**, *48*, 4330–4342. [CrossRef]
5. Jin, Y.; Sharifi, A.; Li, Z.; Chen, S.; Zeng, S.; Zhao, S. Carbon emission prediction models: A review. *Sci. Total Environ.* **2024**, *927*, 172319. [CrossRef]
6. Jiang, Y.; Zhang, J.; Zhang, T.; Liu, X.; Xue, Z.; Liu, Y.; Shi, J. Dynamic carbon emission responsibility factor of electricity. *Proc. CSEE* **2024**, *44*, 7024–7039. [CrossRef]
7. Zhang, N.; Li, Y.; Huang, J.; Li, Y.; Du, E.; Li, M.; Liu, Y.; Kang, C. Carbon measurement method and Carbon meter system for whole chain of power system. *Autom. Electr. Power Syst.* **2023**, *47*, 2–12.
8. Kang, C.; Zhou, T.; Chen, Q.; Xu, Q.; Xia, Q.; Ji, Z. Carbon emission flow in networks. *Sci. Rep.* **2012**, *2*, 479. [CrossRef]
9. Peng, T.; Cai, X.; Qu, Z.; Tang, A.; Shen, R. Research on calculation model of power supply carbon emission factor in regional power grid considering hierarchical and regional decoupling of large-scale power grid. *Proc. CSEE* **2024**, *44*, 894–905. [CrossRef]
10. Chang, L.; Cao, R.; Pang, L.; Tian, H.; Wang, B. Research on optimal dispatching of power system source load coordination considering customer side carbon responsibility under low carbon background. *Distrib. Util.* **2024**, *41*, 46–55. [CrossRef]
11. Zhou, T.; Kang, C.; Xu, Q.; Chen, Q. Preliminary investigation on a method for carbon emission flow calculation of power system. *Autom. Electr. Power Syst.* **2012**, *36*, 44–49.
12. Zhou, T.; Kang, C.; Xu, Q.; Chen, Q. Preliminary theoretical investigation on power system carbon emission flow. *Autom. Electr. Power Syst.* **2012**, *36*, 38–43.
13. Liu, Y.; Li, Y.; Zhou, C.; Song, J.; Deng, H.; Du, E.; Zhang, N.; Kang, C. Overview of carbon measurement and analysis methods in power systems. *Proc. CSEE* **2024**, *44*, 2220–2236. [CrossRef]
14. Li, Y.; Liu, Y.; Yang, X.; He, W.; Fang, Y.; Du, E.; Zhang, N.; Li, J. Electricity carbon metering method considering electricity transaction information. *Proc. CSEE* **2024**, *44*, 439–451. [CrossRef]
15. Yang, Y.; Pan, F.; Zhong, L.; Zhang, J.; Zhao, J. Study on prediction method of node carbon emission factors in power system based on neural network. *Guangdong Electr. Power* **2023**, *36*, 2–9.
16. Dragomiretskiy, K.; Zosso, D. Variational mode decomposition. *IEEE Trans. Signal Process.* **2014**, *62*, 531–544. [CrossRef]
17. Li, G.; Wu, H.; Yang, H. A multi-factor combination prediction model of carbon emissions based on improved CEEMDAN. *Environ. Sci. Pollut. Res.* **2024**, *31*, 20898–20924. [CrossRef]
18. Sha, G.; Liu, L.; Ma, C.; Cong, X.; Duan, Q.; Wu, Y. Orderly charging strategy for electric vehicles considering the vehicle-network interaction. *Distrib. Util.* **2023**, *40*, 46–54. [CrossRef]
19. Chen, J.; Wang, J.; Chen, S.; Zhao, W. Multi-objective optimal scheduling of electric vehicle charging considering carbon emission reduction. *Coal Econ. Res.* **2024**, *44*, 111–122. [CrossRef]
20. Duque, R.; Arbelaez, A.; Díaz, J.F. Online over time processing of combinatorial problems. *Constraints* **2018**, *23*, 310–334. [CrossRef]

21. Schwenzer, M.; Ay, M.; Bergs, T.; Abel, D. Review on model predictive control: An engineering perspective. *Int. J. Adv. Manuf. Technol.* **2021**, *117*, 1327–1349. [CrossRef]
22. He, C.; Geng, J.; Liu, J.; Xu, P.; Zhu, K.; Zhang, J. Multi-period load modeling and regulation potential assessment for electric buses based on data-driven and model-driven. *Power Syst. Prot. Control* **2024**, *52*, 25–34. [CrossRef]
23. Bai, J.; Nguyen-Xuan, H.; Atroshchenko, E.; Kosec, G.; Wang, L.; Wahab, M.A. Blood-sucking leech optimizer. *Adv. Eng. Softw.* **2024**, *195*, 103696. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.