World Electric Vehicle Journal

MDPI

Article

A Two-Stage MILP-GRU-Based Energy Management Framework for Cost-Optimized Solar-Powered EV Charging in Smart Parking Lots

Tallataf Rasheed [1], Abdul Rauf Bhatti [1,*], Muhammad Farhan [1], Ahmed Ali [2] and Akhtar Rasool [3,*]

[1] Department of Electrical Engineering and Technology, Government College University Faisalabad, Faisalabad 38000, Punjab, Pakistan
[2] Department of Electrical and Electronic Engineering Technology, Faculty of Engineering and the Built Environment, University of Johannesburg, Johannesburg 2092, South Africa; aali@uj.ac.za
[3] Department of Electrical Engineering, University of Botswana, Gaborone UB0061, Botswana
* Correspondence: bhatti_abdulrauf@gcuf.edu.pk (A.R.B.); drarasool83@gmail.com or rasoola@ub.ac.bw (A.R.)

Abstract

A transition towards sustainable transportation requires efficient integration of electric vehicles (EVs) with renewable energy sources. This work proposes a two-stage Parking Lot Energy Management Scheme (PLEMS) to minimize charging costs while maximizing solar photovoltaic utilization in commercial parking facilities. In the first stage, the optimization phase is formulated using a mixed-integer linear programming (MILP) that minimizes the overall cost of EV charging while ensuring maximum utilization of locally available PV energy. In the second stage, a gated recurrent unit (GRU)-based deep learning model performs state of charge (SOC) forecasting for EVs parked in the parking lot. Using the predicted SOC for the next time step, the system decides whether each EV will be charged or discharged, ensuring consistency with the cost-optimal MILP strategy from the first stage. The proposed PLEMS achieves up to 62% daily cost savings in charging compared to uncoordinated direct grid charging. However, this cost saving is the outcome of proposed optimization as well as the integration of PV panels in power grid. When compared with nine similar vehicles to grid (V2G)-enabled approaches from the literature, which report cost savings ranging from 9.73% to 52%, the proposed framework shows an improvement of 10% to 52% over these methods. This hybrid MILP-GRU framework offers practical V2G operation and high scalability for large EV fleets in solar-powered smart parking lots.

Keywords: electric vehicles; solar energy; deep learning; mixed-integer linear programming; parking lot energy management scheme

Check for updates

Academic Editor: Michael Fowler

Received: 16 June 2026
Revised: 23 July 2026
Accepted: 26 July 2026
Published: 21 August 2026

Copyright: © 2026 by the authors. Published by MDPI on behalf of the World Electric Vehicle Association. 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

Renewable energy, specifically solar energy, is considered as a primary source of energy nowadays [1]. The rapid rise of electric vehicles (EVs) has driven significant research interest in developing intelligent and sustainable energy management strategies that cater to their unique charging requirements while supporting the integration of renewable energy sources [2,3]. As EV adoption accelerates globally, efficient scheduling algorithms and predictive models become essential for ensuring grid stability, mitigating congestion, and minimizing energy costs [4]. Recent studies also emphasize that scalable grid-friendly control strategies are crucial in managing demand peaks as infrastructure expands, particularly in regions with high EV penetration [5]. In addition, evidence from recent modeling

studies highlights that scalable grid-friendly control strategies are essential for managing demand peaks as the EV adoption increases [6].

Traditionally, optimization techniques such as MILP have played a pivotal role in addressing the complexities of EV charging scheduling [7]. MILP models are highly effective in handling the discrete and continuous variables involved, allowing for the formulation of cost minimization, power balance, and system constraint considerations within a unified framework [8]. For instance, power system optimization studies have demonstrated that incorporating renewable energy forecasts into MILP models can improve the utilization of clean energy while reducing operational costs [9]. Such approaches also extend to managing the interactions between EV charging stations and grid demands, balancing supply and load efficiently [10].

However, the accuracy of these optimization models heavily depends on reliable state estimations, particularly the SOC of EV batteries [11]. Accurate SOC forecasts are necessary for ensuring safe, efficient, and optimized charging schedules [12]. Conventional model-based SOC estimation methods may not fully capture the temporal dependencies and uncertainties present in real-world scenarios. Recent advancements in machine learning have addressed this gap by employing data-driven models such as Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM), and Gated Recurrent Units (GRUs) for SOC prediction [13–15]. Among these, GRU models stand out for their ability to effectively capture long-term dependencies with reduced computational complexity, making them suitable for real-time applications [16].

The integration of machine learning-based SOC prediction with optimization frameworks has gained considerable attention [17,18]. Similarly, the EV load forecasting is also necessary for effective optimization and scheduling of EV charging systems [19]. This symbiosis enables adaptive scheduling that accounts for anticipated battery states, renewable energy output, and dynamic electricity pricing [20,21]. Similarly, Dolgintseva et al. [22] demonstrated multi-step-ahead time series prediction techniques using multiple-output machine learning models, further emphasizing the importance of accurate forecasting in energy systems.

Building on these developments, recent studies have focused on developing cohesive systems that combine SOC prediction with optimal scheduling to improve operational efficiency [23]. By forecasting battery states using GRU or LSTM models, the system can adapt to fluctuations in renewable energy generation and electricity prices, thereby optimizing charging and discharging schedules in a dynamic environment [24]. Such integrated frameworks promise to reduce energy costs, improve resource utilization, and enhance the stability and sustainability of EV charging infrastructure.

This research proposes a comprehensive system that couples MILP-based optimization with GRU-driven SOC forecasting. The approach leverages the strengths of precise system modeling and high-fidelity, real-time predictions to support scalable, resilient, and cost-effective EV charging management. The combined framework not only facilitates better handling of uncertainties but also aligns with the Sustainable Development Goals (SDGs) by promoting clean energy utilization and reducing operational costs.

In summary, the following objectives have been achieved through this research:

- Developing and validating a scalable, decentralized energy management framework for EV charging systems.
- Evaluating the performance of the proposed system for reducing power consumption, operational costs, and peak demand.
- Utilizing renewable energy forecasts and actual EV charging data for load forecasting to optimize scheduling.

By achieving these objectives, this work has contributed to the broader effort of creating resilient, sustainable, and efficient energy infrastructure capable of supporting the evolving landscape of EV integration and renewable energy deployment.

The primary contribution of this work lies in the development of a practical two-stage hybrid PLEMS framework that integrates MILP-based scheduling with a GRU-based surrogate learning model for scalable EV fleet management. Rather than introducing a new learning paradigm, the proposed approach applies supervised approximation of MILP-derived scheduling trajectories to enable efficient decentralized execution. This system-level integration supports cost-effective scheduling and improved solar energy utilization in parking lot environments.

2. Methodology

The scheduling of electric vehicle (EV) charging has been widely investigated in the literature, with most existing approaches relying on centralized control algorithms to reduce cost and optimize charging power [25]. While effective for small-scale systems, centralized methods face scalability issues when a large number of EVs are integrated, as their aggregated demand can significantly alter the grid load profile [26]. This necessitates the development of more robust solutions that ensure cost-effectiveness while maintaining grid stability [27].

Decentralized control has emerged as a promising alternative due to its ability to manage a large fleet of EVs autonomously without imposing excessive computational burden on a central server [28]. Unlike centralized approaches, decentralized systems allow each EV or charging station controller to make local decisions based on limited but relevant information [29]. This autonomy enhances scalability, reduces communication overhead, and improves robustness against failures [30]. However, ensuring that local decisions collectively satisfy global objectives such as minimizing overall charging cost and preserving grid reliability remains a significant challenge [31]. Furthermore, accurate scheduling requires detailed knowledge of EV-specific parameters such as arrival times, departure times, and future base load, which complicates techno-economic analysis [32].

To address these challenges, this work proposes a decentralized scheduling framework, termed the Parking Lot Energy Management Scheme (PLEMS), which coordinates EV charging and discharging while considering renewable integration, grid interaction, and user requirements. The two-stage PLEMS framework is presented as follows: Section 2.1 formulates the offline MILP optimization; Section 2.2 introduces the GRU-based scheduling model and its architecture; Section 2.3 details the dataset and forecasting procedure used to train it; and Section 2.4 describes its online deployment. The model architecture is presented ahead of its training data for conceptual clarity, consistent with common practice in machine learning methodology sections.

The proposed PLEMS operates in two sequential phases, as illustrated in Figure 1. In the optimization phase, forecasts of EV load demand and PV generation are utilized to determine the optimal charging cost and SOC trajectories of EVs within the parking facility. In the subsequent scheduling phase, an energy management scheme is applied to refine the optimization results, dynamically deciding when EVs should charge, discharge, or remain idle. This two-phase framework ensures both cost minimization and efficient utilization of available energy resources.

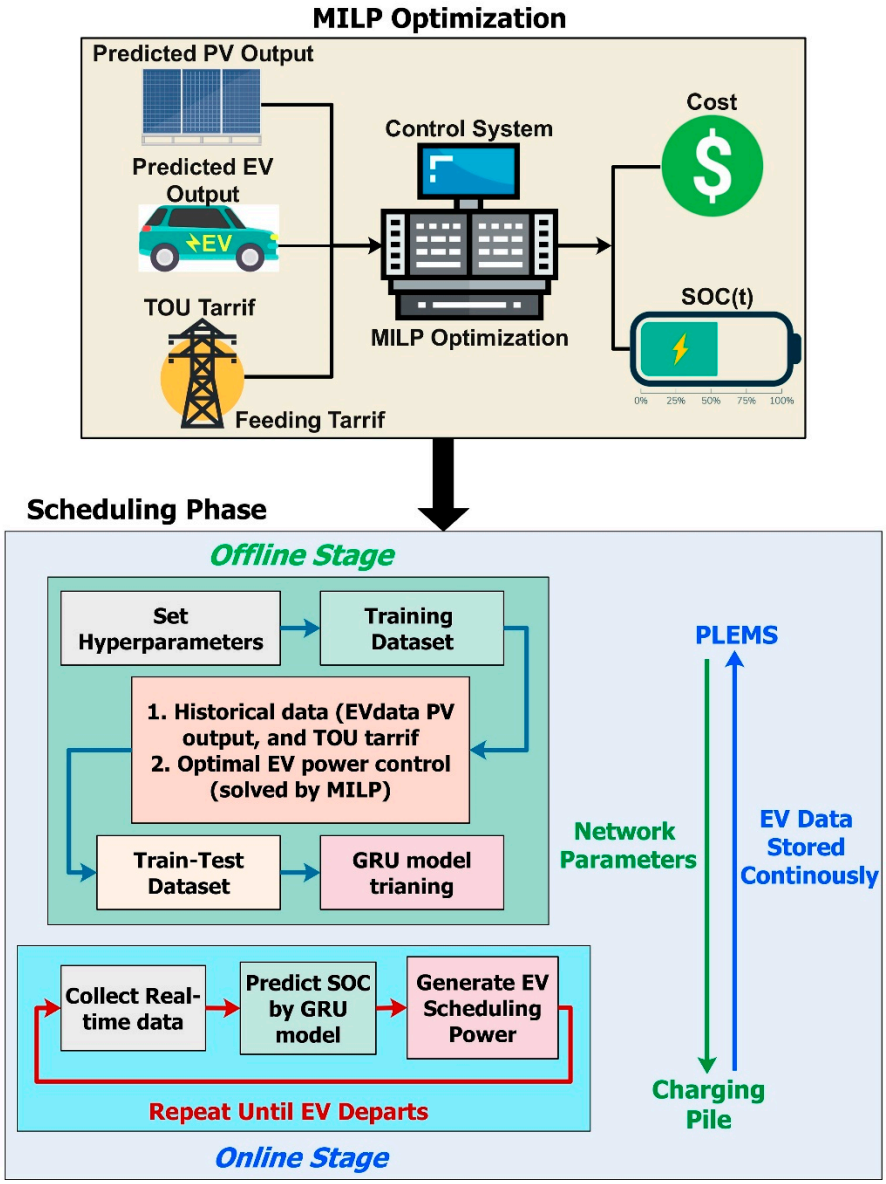

Figure 1. The Parking Lot Energy Management Scheme.

2.1. Optimization Phase

In the proposed PLEMS framework, the optimization phase is formulated as a MILP problem to minimize the overall cost of EV charging while ensuring effective utilization of locally available PV energy, as shown in Figure 1. The optimization integrates forecasted EV load demand, PV generation, time-of-use (TOU) tariffs, and feed-in tariffs with the behavioral and technical parameters of EVs. Historical EV load and PV generation profiles from an EV charging station in Boulder City, NV, USA (described in detail in Section 2.3) are used to represent realistic operating conditions and to provide training data for the subsequent learning phase. The MILP does not assume perfect knowledge of future conditions; instead, it is solved using day-ahead forecasts of EV load demand, PV generation, and time-of-use tariffs, generated via the forecasting framework described in Section 2.3 using [33]. By optimizing offline on representative daily scenarios under these realistic forecast inputs, the MILP produces forecast-optimal SOC trajectories and charging/discharging decisions that respect all physical and operational constraints. These trajectories then serve as high-quality, constraint-satisfying ground-truth labels for training the GRU model in the scheduling phase.

Building on the forecasted inputs described above, the MILP formulation begins with the fundamental power balance equation at the point of common coupling (PCC):

$$P_g(t) = F_d(t) - P_s(t) + \sum_{i=1}^{N} P_{Evi}(t) \qquad \forall\, t \in T \tag{1}$$

where

- $P_g(t)$: net power exchanged with the grid at time t (positive = import, negative = export);
- $F_d(t)$: forecasted load demand at time t;
- $P_s(t)$: forecasted PV generation at time t;
- $P_{Evi}(t)$: charging/discharging power of EV i at time t;
- N: total number of EVs;
- T: scheduling horizon (discretized into 15 min intervals).

The MILP guarantees a global optimum with respect to the objective function as formulated; however, this objective function is evaluated using day-ahead forecasts of EV demand, PV generation, and tariff signals rather than their realized values. The resulting solution is therefore globally optimal conditional on forecast accuracy; i.e., it is the best possible schedule given the information available at decision time, not a guarantee of optimality against the true, realized operating conditions.

The optimization minimizes the total energy cost, defined as

$$\min\,(P_{EViC},\, P_{EViD}) \sum_{t=1}^{T} (T_u(t) \times DF_{im}(t) \times F_d(t)) + (C_f(t) \times DF_{ex}(t))(1 - F_d(t)) \tag{2}$$

where

- $T_u(t)$: Time of Use tariff at time t;
- $C_f(t)$: feed-in tariff at time t;
- $DF_{im}(t)$: binary variable, equal to 1 if grid import occurs at time t;
- $DF_{ex}(t)$: binary variable, equal to 1 if grid export occurs at time t.

Equations (3) and (4) define the threshold-based rules governing whether the PLEMS should operate in import or export mode:

$$DF_{im}(t) = \begin{cases} 1, & if\ F_d(t) \geq b_u, \\ 0, & otherwise, \end{cases} \tag{3}$$

$$DF_{ex}(t) = \begin{cases} 1, & if\ F_d(t) \leq b_l, \\ 0, & otherwise, \end{cases} \tag{4}$$

with

- b_u: upper threshold (set to 70% of forecasted demand, meaning imports are enforced when demand exceeds this level);
- b_l: lower threshold (set to 30%, meaning exports occur when demand is below this level).

To prevent simultaneous grid import and export, the exclusivity constraint is imposed, given in Equation (5). Unlike conventional MILP formulations for prosumer EV charging that rely primarily on price signals (TOU vs. feed-in tariffs) with basic exclusivity to guide grid interactions [24], the proposed optimization incorporates proactive threshold-based rules (upper threshold of 70% and lower threshold of 30% of forecasted demand) along with strict exclusivity and bounded import/export capacities. These mechanisms aggressively maximize solar self-consumption and promote grid-friendly load flattening even when tariff differences are modest, yielding higher-quality (lower-cost) optimal SOC

trajectories as training labels for the subsequent scheduling phase. It should be noted that the thresholds defined in Equations (3) and (4) regulate grid interaction based on demand conditions. In addition, a separate SOC-based operational constraint is imposed to ensure user mobility and prevent excessive battery discharge. Specifically, bidirectional discharge (V2G) is permitted only when the battery SOC exceeds 70%, while charging is prioritized when SOC falls below 30%. This additional constraint prevents profit-driven deep discharge while preserving minimum departure readiness, ensuring that the optimization remains cost-minimizing within a user-service-constrained feasible region. The 70/30 demand thresholds regulate grid interaction behavior under varying load conditions, while SOC-based 70/30 limits independently enforce battery safety and user mobility constraints during V2G participation.

$$DF_{im} + DF_{ex} = 1 \tag{5}$$

To ensure that discharge decisions are fully specified and reproducible, bidirectional discharge for EV i at time t is permitted if and only if all of the following hold:

(i) SOC exceeds the safety threshold, $SOC_i(t) > 70\%$;

(ii) The vehicle is not committed to imminent departure within the current control horizon (idle status);

(iii) Grid demand is in the export-favorable regime, $F_e(t) \le b_1$ (Equation (4));

(iv) Discharging at time t does not jeopardize the departure-readiness constraint; i.e., the projected SOC at departure remains $SOC_i\left(t_{dep}\right) \ge 85\%$.

If any of (i)–(iv) are not satisfied, EV i defaults to charging or idle mode, never discharge. This can be expressed as Equation (6).

$$V2G_i(t) = 1 \leftrightarrow [SOC_i(t) > 0.70] \wedge [Idle_i(t)] \wedge [F_e(t) \le b_l] \wedge [SOC_i(t_{dep}) \ge 0.85] \tag{6}$$

Consequently, "limited bidirectional capability" in this work refers specifically to discharge events that satisfy Equation (6), a bounded subset of otherwise-idle, high SOC, export-favorable intervals rather than continuous, on-demand bidirectional operation as required under full V2G.

The power balance from Equation (1) is re-expressed as a constraint:

$$P_g(t) = F_d(t) - P_s(t) + \sum_{i=1}^{N} P_{Evi}(t) \qquad \forall\, t \in T \tag{7}$$

The import/export capacities are limited as

$$B_{ex}(t) \cdot b_l \le P_g^{ex}(t) \le 0 \qquad \forall t \in T \tag{8}$$

$$0 \le P_g^{im}(t) \le B_{im}(t) \cdot b_u \qquad \forall t \in T \tag{9}$$

where $P_g^{im}(t)$ and $P_g^{ex}(t)$ denote the import and export components of $P_g(t)$, and $B_{im}(t)$, $B_{ex}(t)$ are participation binaries for import and export.

The EV power is decomposed into charging and discharging contributions:

$$P_{EV,i}(t) = P_{vi}^{C}(t) + P_{vi}^{D}(t), \qquad \forall t \in T,\ \forall i \in N \tag{10}$$

with bounds

$$-B_{vi}^{D}(t) \cdot P_{vi}^{D,u}(t) \le P_{vi}^{D}(t) \le 0, \quad \forall t \in T,\ \forall i \in N \tag{11}$$

$$0 \le P_{vi}^{C}(t) \le B_{vi}^{C}(t) \cdot P_{vi}^{C,u}(t), \qquad \forall t \in T,\ \forall i \in N \tag{12}$$

$$B_{vi}^{C}(t) + B_{vi}^{D}(t) = 1, \qquad \forall t \in T,\ \forall i \in \mathrm{N} \tag{13}$$

where

- $P_{vi}^{C}(t)$, $P_{vi}^{D}(t)$: charging and discharging power of EV i;
- $P_{vi}^{C,u}(t)$, $P_{vi}^{D,u}(t)$: maximum charging and discharging power of EV i;
- $B_{vi}^{C}(t)$, $B_{vi}^{D}(t)$: binary participation variables for charging and discharging, ensuring mutual exclusivity.

The SOC dynamics of each EV battery are modeled from [24]:

$$S_i(t) = S_i(t-1) + \left(\eta_i^C P_{vi}^C(t) + \frac{P_{vi}^D(t)}{n_i^D}\right)\frac{\tau}{C_i} \qquad \forall t \in T,\ \forall i \in N \tag{14}$$

- $S_i(t)$: SOC of EV i at time t;
- η_i^C, n_i^D: charging and discharging efficiencies;
- C_i: battery capacity of EV i (30 kWh assumed);
- τ: time step (15 min in this work).

Equation (14) ensures that the SOC is dynamically updated at each interval, accounting for both charging and discharging activities while reflecting conversion losses. SOC bounds ($S_i^{min} \le S_i(t) \le S_i^{max}$) and departure requirements ($S_i(D_i) \ge S_i^{dep}$) are imposed to guarantee that user mobility needs are met. This formula calculates how the battery's SOC changes from one time step to the next. It takes the current SOC, adds energy when charging (divided by charging efficiency), or subtracts energy when discharging (multiplied by discharging efficiency). It makes sure charging and discharging do not happen at the same time and keeps the SOC between ($S_i^{min} \le S_i(t) \le S_i^{max}$).

The MILP is solved on a day-ahead basis, producing cost-optimal charging and discharging schedules. These solutions not only minimize charging costs while maximizing renewable utilization but also generate labeled data for training deep learning models in the scheduling phase. The GRU model, trained on MILP-derived optimal schedules, approximates these solutions in real time under uncertain operating conditions such as variable PV generation and stochastic EV arrivals and departures.

2.2. Scheduling Phase

The second phase of the proposed PLEMS framework focuses on scheduling, where the charging and discharging actions of individual EVs are coordinated in real time. Initially, the scheduling procedure is derived for a single EV and then applied to a fleet of multiple EVs. The decision relies on predicting the SOC at the subsequent time step $(t+1)$, which is obtained using a gated recurrent unit (GRU)-based neural network. The choice of GRU is motivated by previous work, which demonstrated its superior performance compared to other deep learning models such as LSTM and RNN for time-series load prediction [33]. Unlike a standard recurrent unit, the GRU regulates information flow through two gating mechanisms: an update gate and a reset gate. The update gate determines how much of the previous hidden state is retained versus replaced by new information, while the reset gate determines how much past information to discard when computing a candidate hidden state. This gating structure allows the GRU to capture temporal dependencies in SOC evolution across time steps while mitigating the vanishing gradient problem associated with simple RNNs, and with fewer parameters than an LSTM, resulting in faster training and lower computational overhead, an important consideration given the frequent retraining and large fleet sizes (n EVs) targeted by this work. A comparative analysis conducted in this work confirmed that GRU provides the highest accuracy in SOC forecasting, thereby making it the most suitable candidate for efficient scheduling.

During offline training, the GRU learns the optimal SOC trajectory generated by the MILP; however, in real-time operation the model uses the actual measured SOC of each

connected EV at every interval and predicts the next-step desired SOC, which is then translated into charging/discharging power using Equation (14), ensuring closed-loop physical consistency rather than open-loop policy imitation.

At each time step t, the scheduling function determines the EV charging or discharging power using the SOC difference between two consecutive time steps. This relationship is given in [24] as

$$P_{V_i}(t) = \begin{cases} \dfrac{S_i(t+1) - S_i(t)}{\tau \eta_{ci}} C_i & S_i(t+1) - S_i(t) \geq 0 \\ \dfrac{S_i(t+1) - S_i(t)}{\tau / \eta_{di}} C_i & S_i(t+1) - S_i(t) < 0 \end{cases} \quad (15)$$

where

- $P_{V_i}(t)$: charging/discharging power of EV i at time t;
- $S_i(t)$, $S_i(t+1)$: SOC of EV i at time t steps and $t+1$;
- η_{ci}, η_{di}: charging and discharging efficiencies;
- C_i: battery capacity of EV i;
- τ: scheduling interval (15 min).

In the above formulation, $S_i(t)$ denotes the SOC of the ith electric vehicle battery at time interval t, expressed as a fraction of its total energy capacity. The parameters η_{ci} and η_{di} correspond to the charging and discharging efficiencies, capturing conversion losses during energy transfer between the grid and the battery. Binary variables are employed to ensure mutually exclusive charging and discharging operations at each time step, preventing physically infeasible simultaneous actions. The electricity tariff parameters reflect time-of-use (TOU) pricing, allowing the optimization framework to shift charging demand toward low-cost periods or locally available PV generation. Collectively, these terms ensure that the scheduling decisions remain physically feasible, economically optimal, and consistent with real-world EV charging behavior.

Equation (15) works in reverse as compared to Equation (13). The GRU predicts the "desired" SOC for the next time step. The system then calculates how much power is needed right now to reach that desired SOC from the current (actual measured) SOC. It uses the right efficiency (charging or discharging) depending on whether the battery needs to gain or lose energy. The power is limited to what the charger can handle and adjusted if it would push the SOC outside safe limits.

- Equation (14) is for planning the perfect schedule ahead of time (offline).
- Equation (15) is for deciding the actual power in real-time based on the GRU's prediction.

Both follow the same basic physics of how batteries charge/discharge, so they stay consistent. The real-time version also uses the latest measured SOC to correct any small prediction mistakes and keep everything safe and stable. Equation (15) ensures that the PLEMS schedules charging power proportionally to the SOC variation while accounting for efficiency losses and battery capacity. Consequently, SOC forecasting becomes the central task of the scheduling phase. Given its temporal dependence, SOC prediction is naturally formulated as a time-series problem, where one-step-ahead predictions are iteratively propagated over the full scheduling horizon. The overall scheduling procedure is illustrated in Figure 1.

The GRU is embedded within the PLEMS to forecast SOC and enable decentralized scheduling across all connected EVs. The structure is divided into two stages: an offline training stage at the PLEMS level and an online execution stage at each charging pile. The proposed two-stage framework tightly couples physics-based optimization with data-driven prediction. In the first stage, a MILP optimizer is solved over the entire prediction

horizon using day-ahead forecasts of PV generation and vehicle availability (Section 2.3). Electricity tariffs and feed-in tariffs are treated as known in advance; in this work, the tariff structure is adopted directly from [24] to ensure a consistent basis for comparison with the reproduced baseline. This yields the globally optimal SOC trajectory that minimizes total energy cost while respecting all battery, power, and operational constraints as given in Equation (3) to Equation (13).

This MILP-derived optimal SOC profile is then used as the ground-truth supervisory label for training the proposed GRU models. Importantly, these supervisory labels are generated from offline MILP solutions under realistic day-ahead forecasts taken using previous work [33] (not perfect hindsight), ensuring the GRU learns a robust approximation of optimal decisions that remains effective even when online forecasts deviate slightly from historical patterns. Specifically, at each time step t the input feature vector containing four features: (i) the current SOC $S(t)$, (ii) PV output $P_s(t)$, (iii) forecasted load demand $F_d(t)$, and (iv) the TOU tariff $T_u(t)$. Importantly, these labels are generated using day-ahead forecasts rather than perfect hindsight, ensuring the GRU learns a robust approximation of optimal decisions even under real-world uncertainties like variable PV output or stochastic EV arrivals. The output layer contains a single node representing the predicted desired SOC at the next time step (the optimal target value from MILP). This input is paired with the MILP-computed optimal $SOC(t)$ as the target value. By training the GRU on MILP-derived optimal SOC trajectories over representative daily scenarios, the network learns the functional relationship between system inputs and the corresponding cost-optimal SOC targets. This mapping enables real-time execution without repeatedly solving the full MILP, thereby supporting fast and scalable online scheduling under uncertain operating conditions. During inference, the trained GRU directly predicts the next-step desired target SOC.

2.3. Offline Training Stage

The GRU network is trained to predict the desired next-step SOC (i.e., the optimal target SOC that the MILP optimizer would achieve at the next 15 min interval under optimal conditions). The GRU approximates this optimal target trajectory, enabling real-time scheduling decisions via Equation (14) that correct deviations using the measured current SOC. The dataset covers 11 months of charging activities (January–November 2023) collected from the Boulder City charging station in the USA. The dataset consists of real historical charging transaction records from city-owned EV charging stations in Boulder, Colorado, USA (sourced from the City of Boulder's open data portal for the period January–November 2023). This includes actual transaction start/end times, energy delivered (kWh), charging duration, and station locations, capturing natural variability in EV arrival/departure patterns and daily parking behavior. PV generation data for the same period were obtained from the City of Boulder's Generation Solar project. Specifically, data came from the real-time monitoring system at the 1500 Pearl Parking Garage solar installation. This was queried over the same January–November 2023 date range. Both the EV charging transaction data and the PV generation data are therefore real, measured observations, not simulated values.

In contrast, several EV battery parameters are modeling assumptions, not observed quantities. The source datasets do not record per-vehicle SOC. The initial SOC (0.6) and departure SOC threshold (0.85) are fixed values. They are applied uniformly across all simulated EVs. Battery capacity (30 kWh) is also an assumed representative value (Section 2.1). The real-time SOC trajectory during simulation is not observed either. It is computed entirely by the model's own charging/discharging dynamics (Equation (14)), starting from the assumed initial SOC.

Day-ahead forecasts for both EV load demand $F_d(t)$ and PV generation $P_s(t)$ are generated using the Improved Supervised Learning (ISL) deep-learning forecasting framework from our earlier work [33]. This framework reframes each time series as a supervised learning problem, using lag-based features (at $t-2$, $t-1$) with a GRU model. It does not use raw historical values directly as inputs. EV load forecasting at the 15 min resolution used in the work [33] reports a normalized MAE of 0.014 and a normalized RMSE of 0.213 on held-out test data. The same ISL methodology was applied to generate the PV generation forecasts used here, which confirms that $F_d(t)$ and $P_s(t)$ carry realistic, non-trivial forecast error relative to the actual realized demand and generation. They do not simply reproduce the historical ground truth. These forecasts simulate practical aggregator-level predictions available in real operations (e.g., via weather services for PV and historical occupancy for EV arrivals), providing realistic inputs for the MILP optimizer and GRU evaluation. The chronological train–test split ensures the held-out test set reflects unseen but representative real-world days with inherent stochasticity (e.g., variable weather impacting PV or fluctuating EV numbers). Figure 2a,b present the predicted EV charging demand and PV generation using the method in [33], while Figure 2c illustrates the daily EV charging profile, and Figure 2d shows the corresponding TOU tariff structure considered from [24] for consistent benchmarking with the reproduced baseline. The chronological train–test split ensures the held-out test set reflects unseen but representative real-world days with inherent stochasticity (e.g., variable weather impacting PV or fluctuating EV numbers), captured without synthetic disturbances. The use of these realistic, independently forecasted day-ahead inputs (rather than perfect actual values) for generating MILP optimal labels ensures that training inputs incorporate natural prediction uncertainties similar to those encountered in online deployment. This alignment prevents the GRU from learning an overly idealized policy and promotes robustness to real-world forecast errors.

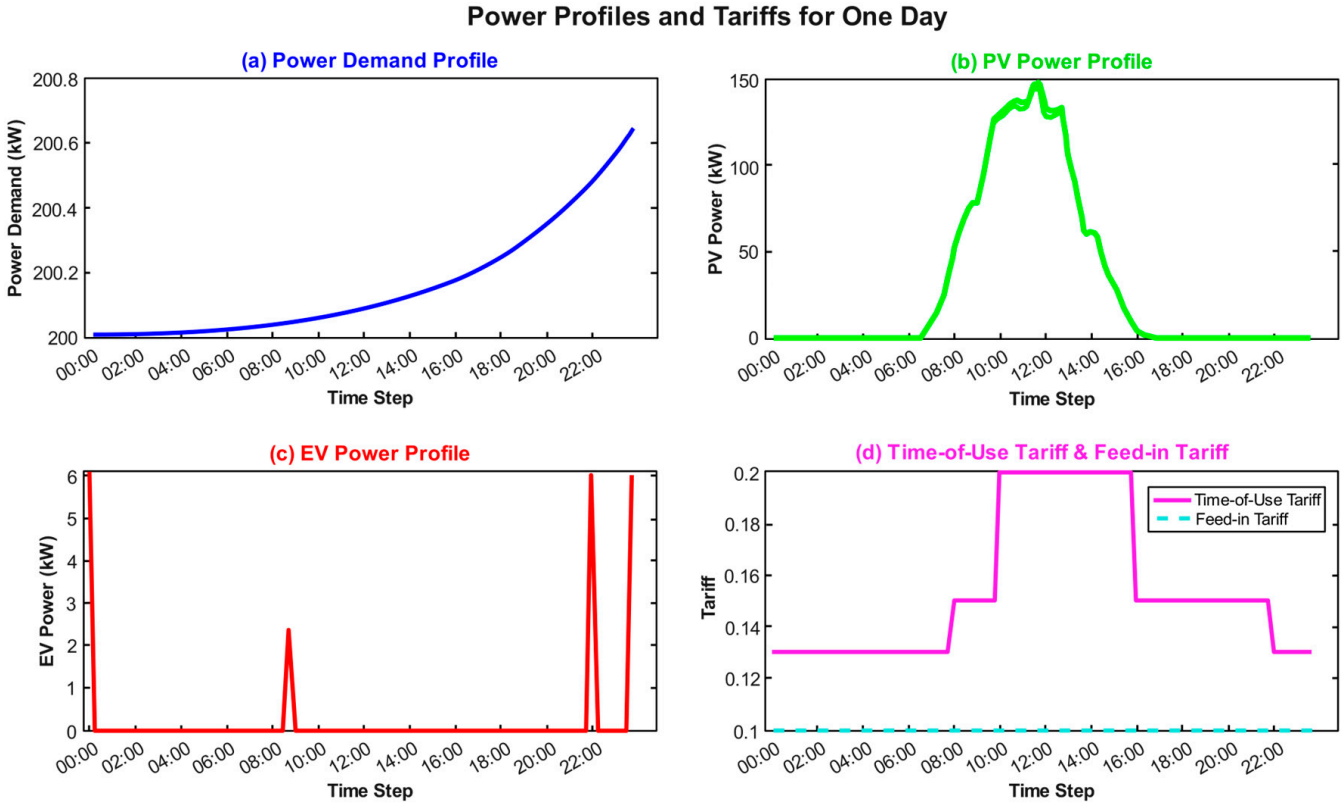

Figure 2. Overview of the 1500 Pearl Street Parking Garage in Boulder, USA, showcasing (**a**) Predicted EV Load, (**b**) Predicted PV output, (**c**) Daily EV profile at the site, and (**d**) Electricity tariffs at different time steps.

The dataset was divided into training and test sets using a 70–30% chronological split to preserve temporal order and prevent data leakage. Specifically, the first 70% of samples were used for training, while the remaining 30% were reserved as an independent test set for final evaluation. Input features were constructed using a supervised learning framework as in [33], and both training and test inputs were reshaped into 3D format [samples, timesteps = 1, features] as required by recurrent layers. All input features were standardized using Standard Scaler fitted only on the training set, and the target variable (SOC) was normalized to the range [0, 1] using Min Max Scaler for numerical stability.

Each model consists of three stacked recurrent layers (128 → 64 → 32 units) with Batch Normalization and Dropout (0.3–0.4) applied after each layer, followed by two dense layers (32 and 16 units, ReLU activation) and a linear output layer, as shown in Figure 3. All models were compiled using the Adam optimizer with an initial learning rate of 0.0005 and mean squared error (MSE) as the loss function. Models were trained for a maximum of 300 epochs with a batch size of 32. Model performance was evaluated using mean absolute error (MAE), root mean squared error (RMSE), and coefficient of determination (R^2) computed on both training and test sets.

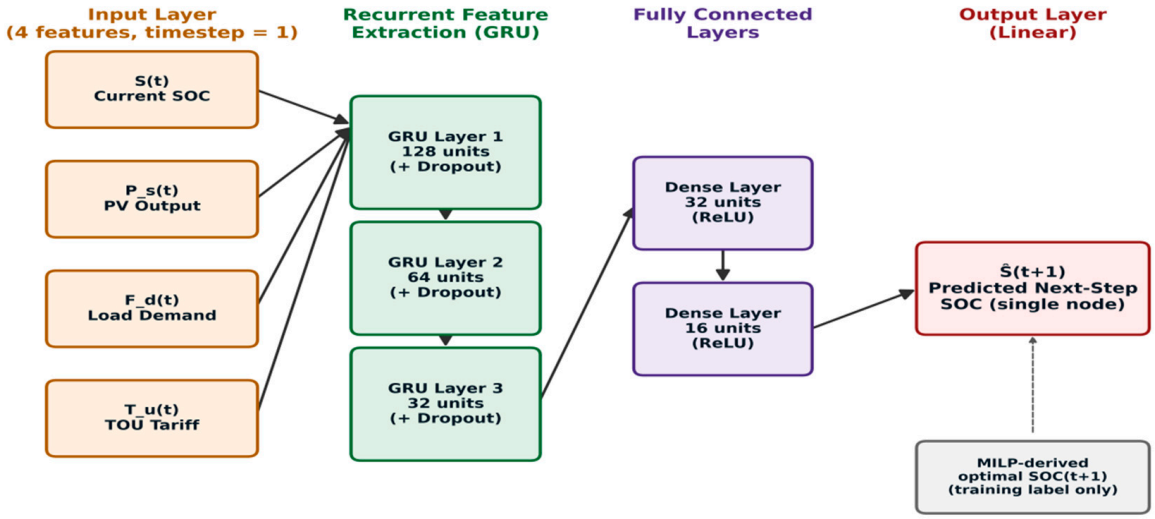

Figure 3. Input and output structure of the GRU-based SOC prediction model.

The GRU architecture consisting of three layers (128 → 64 → 32 units) was selected based on a balance between predictive performance and computational efficiency. The progressively decreasing layer size enables effective feature extraction while limiting model complexity, which is important for real-time deployment at charging-pile level. Preliminary experiments with larger architectures showed only marginal accuracy improvement but increased training time and overfitting risk. Hyperparameters such as layer size, dropout rate, and learning rate were tuned empirically using validation performance rather than exhaustive grid search. Overfitting was mitigated through dropout regularization, chronological train–test splitting to avoid temporal leakage, and normalization of input and target variables, ensuring stable convergence and good generalization on unseen data.

2.4. Online Execution Stage

Once trained, the GRU parameters are deployed at each charging pile. When an EV arrives, the charging pile reads its current SOC and combines it with the real-time inputs described in Section 2.2, ($P_s(t)$, $F_d(t)$, $T_u(t)$), obtained from the PLEMS. The GRU then predicts the SOC for the next interval, which is used in Equation (14) to determine the charging or discharging power. The scheduling horizon spans a full day (00:00–24:00), discretized into 96 intervals of 15 min each. Each EV is assumed to start with an SOC of

0.6 and a battery capacity of 30 kWh, with charging/discharging efficiencies fixed at 95%. A departure SOC threshold of 0.85 is imposed to ensure user mobility requirements are met. The MILP-generated optimal schedules serve as reference labels for GRU training, allowing the network to replicate cost-optimal strategies under uncertainty.

The proposed two-stage PLEMS framework is specifically designed to overcome the limitations of standalone MILP in real-world deployment. While pure MILP can compute globally optimal schedules under perfect foresight, it is computationally intensive and requires centralized high-performance resources, making repeated online re-optimization impractical for large fleets (e.g., 100 EVs) or edge devices at charging piles. Moreover, MILP's sensitivity to forecast errors and long solve times hinders fast, decentralized decision-making. By using offline MILP to generate high-quality optimal labels for training a lightweight GRU model, the approach enables real-time, decentralized inference with excellent scalability and robustness to minor uncertainties, advantages not achievable with MILP alone. Here, "near-optimal performance" refers specifically to the GRU's demonstrated ability to inductively generalize from MILP-derived training labels to unseen operating days, quantified in Section 3 by a test RMSE of 0.97, test R^2 of 0.9754, and test MAE of 0.29 relative to the MILP-optimal SOC targets, translating into consistent empirical cost savings of 45–62% on held-out data. As with any inductively trained model, this performance is a statistical property of the test distribution used for evaluation and is not a formal optimality guarantee; it is expected to hold to the extent that future operating conditions remain consistent with the patterns present in the training data. This hybrid design thus provides a practical, efficient solution for solar-integrated EV parking lot management, balancing theoretical optimality with deployment feasibility.

2.5. Scalability and Applications

The proposed framework is validated for 100 EVs, demonstrating its scalability for large parking facilities. The MILP is solved once per day, offline and in advance, using MATLAB R2021b's built-in intlinprog solver on standard laptop class hardware (Intel Core i7-1165G7 @ 2.80 GHz, 16 GB RAM) to generate that day's optimal labels. Solve time for this offline step depends on the joint scheduling problem's scale, which grows with the number of EVs and the length of the scheduling horizon, since each additional EV and time interval introduces further charge/discharge and grid-interaction decisions into the branch-and-bound search. This is perfectly acceptable, however, even if it takes minutes or hours, as it is a one-time, offline pre-computation performed once per day. It is never executed within the real-time control loop and therefore does not affect real-time scalability; online scheduling relies exclusively on the trained GRU, which executes in <1 s per decision. At deployment, each of the 100 EVs runs the same trained GRU independently in <1 s with zero inter-vehicle communication or centralized coordination. Its modular design supports further extension to larger fleets or interconnected zones. The model prioritizes renewable utilization by scheduling EVs to charge during peak PV production, while grid imports are shifted to off-peak hours. Vehicle-to-grid (V2G) participation follows the discharge conditions defined in Equation (6) (SOC > 70%, idle status, export favorable demand, and guaranteed departure readiness), balancing user satisfaction and grid support.

In residential or small-scale contexts, the framework can be simplified by reducing dimensionality (e.g., fewer EVs, basic tariff models) while retaining its optimization logic. Compatibility with existing smart charging infrastructure ensures practical feasibility, while integration with TOU tariffs and demand response programs enhances economic sustainability. Future enhancements may include multi-objective optimization, considering battery degradation, carbon emissions, and user satisfaction alongside cost minimization.

3. Results and Discussion

This work evaluates the proposed PLEMS by comparing its performance against a SOC baseline and one literature-based optimization model [24]. For a fair comparison, the model in [24] is also implemented using our dataset to obtain realistic comparative results. The comparative analysis focuses on cost savings, power consumption, and SOC optimization to highlight improvements in efficiency, affordability, and sustainability.

Figure 4 illustrates the optimal SOC trajectories for six randomly selected EVs, as computed by the day-ahead MILP optimizer. The *X*-axis represents the time interval (15 min resolution, from 0 to 95, covering 24 h), while the *Y*-axis shows the SOC in fraction (0.4 = 40% to 1.0 = 100%). The step plots depict SOC evolution over a 24 h horizon discretized into 96 intervals of 15 min resolution, with each subplot representing one EV. The horizontal segments indicate periods of idle operation, where no charging or discharging occurs, reflecting the optimizer's decision to conserve energy or await favorable conditions. The vertical steps represent transitions at decision points, corresponding to scheduled charging (upward) or limited discharging (downward, if V2G is active), aligned with peak PV generation and low time-of-use (TOU) tariffs to minimize costs while ensuring departure SOC requirements ($\geq$0.85) are met. Changes are physically realistic and gradual, typically 1–4% per active interval at standard Level 2 charging rates, emphasizing the framework's ability to integrate renewable energy effectively while maintaining battery constraints and user mobility needs.

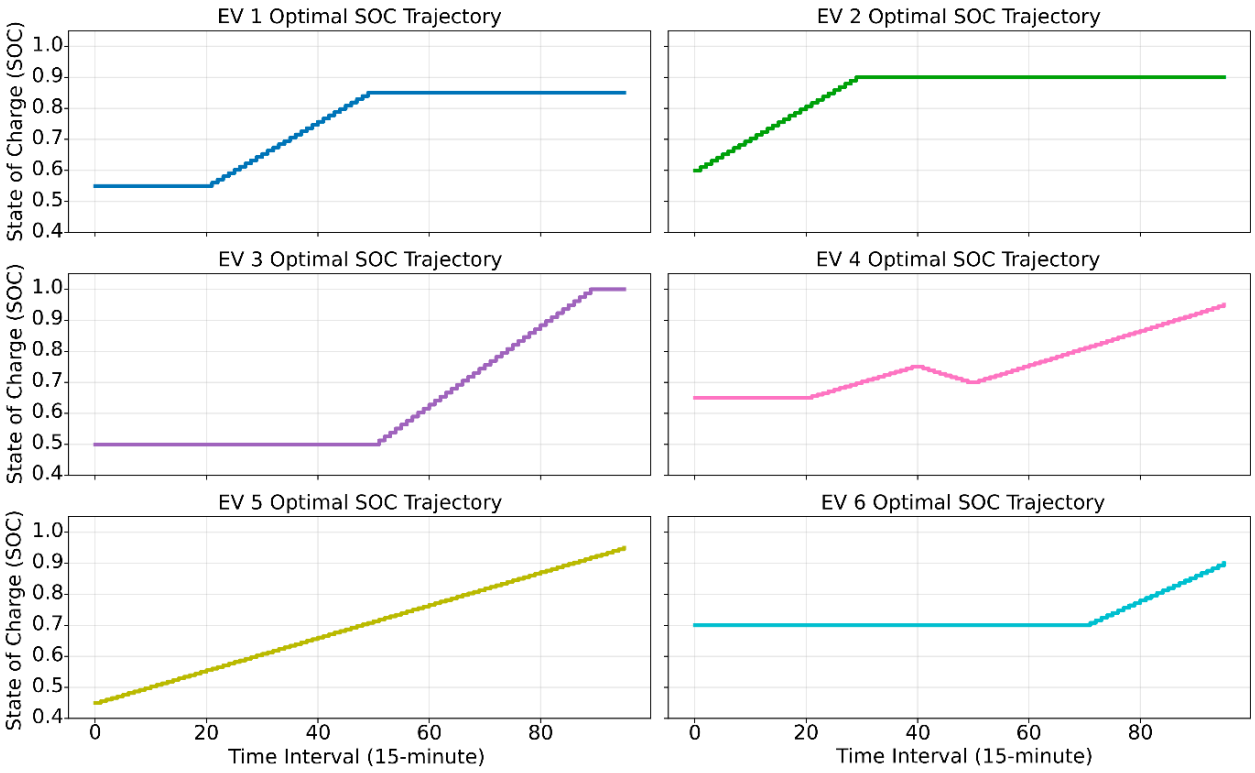

Figure 4. Optimal SOC Trajectories for 6 Selected EVs (MILP Optimization).

The charging cost is shown in Figure 5, where the proposed PLEMS achieves daily cost reductions of 45–62% compared to uncoordinated charging, with a total weekly saving of 51.86% relative to the non-optimized baseline, and a total weekly saving of 9.56% relative to the literature-based benchmark [24]. These results are consistent with, or superior to, savings reported in recent V2G-enabled parking lot studies, which typically range from 20 to 65% against similar uncoordinated baselines [34–36].

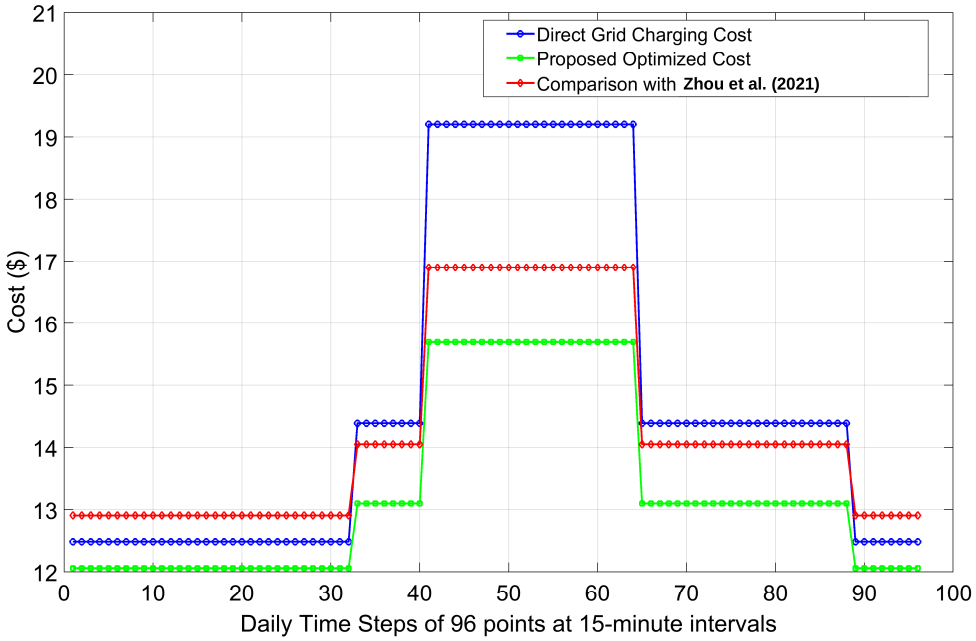

Figure 5. Optimized cost with MILP and proposed objective function for a day [24].

A day-by-day comparison of charging costs is presented in Table 1. The proposed objective function consistently reduces cost relative to both the non-optimized baseline and the literature method [24]. On a per-day basis the proposed method achieves savings versus the non-optimized baseline in the range of 45.68–62.21%, and versus the literature method in the range of 1.71–21.82%. Aggregating the seven days, the total weekly cost under the non-optimized scenario is $17,490, while the proposed method yields a cumulative weekly cost of $8420, a 51.86% reduction in total weekly expenditure. Relative to the literature optimization the total weekly reduction is 9.56%. The mean daily saving (simple arithmetic mean of the seven daily percentage savings vs. non-optimized) is 51.52%, and the mean daily saving versus the literature method is 8.98%. These corrected aggregate figures are consistent with the per-day values reported in Table 1 (Day 1–Day 7).

Table 1. Daily cost comparison.

Days	Direct Grid Charging Cost ($)	Literature-Based Optimized Cost ($) [24]	Proposed Optimized Cost ($)	Cost Saving by Proposed Scheme (%)	
				Compared to Non-Optimized Charging	Compared to Literature-Based Charging [24]
Day 1	2590	1410	1370	47.10	2.84
Day 2	2430	1450	1320	45.68	8.97
Day 3	2310	1300	1100	52.38	15.38
Day 4	3070	1210	1160	62.21	4.13
Day 5	2570	1650	1290	49.80	21.82
Day 6	2230	1170	1150	48.43	1.71
Day 7	2290	1120	1030	55.02	8.04

The non-optimized baseline represents realistic current practice at most charging facilities; EVs charge immediately upon arrival, drawing power directly from the grid without coordination, scheduling, or utilization of on-site PV generation. This scenario is deliberately chosen, not as an artificially weakened comparison, but because it reflects

the actual mode of operation this work aims to improve upon; as EV adoption grows, such uncoordinated, grid-only charging is increasingly unable to scale without risking local transformer and feeder overload. Benchmarking against this realistic baseline is also standard practice in the V2G scheduling literature; as shown in Table 2, all nine comparable studies surveyed in this work similarly report savings relative to uncoordinated charging baselines, ensuring the reported PLEMS improvements are directly comparable to prior reported results in the field.

It should be noted that because the non-optimized baseline does not utilize on-site PV generation, part of this reported cost reduction reflects the value of PV availability itself, in addition to the scheduling optimization performed by PLEMS; isolating these two effects independently is left for future work.

The temporal cost dynamics are further illustrated in Figure 6, which presents the monthly cost trend over the evaluation period. It demonstrates that the proposed approach not only reduces daily expenses but also maintains consistent benefits across longer horizons, with average monthly savings aligning closely with the daily figures (approximately 50–55% vs. uncontrolled charging). To contextualize these results, Table 2 compares the proposed PLEMS with recent V2G-enabled approaches that use uncoordinated charging as a baseline. While full V2G implementations can achieve up to 65% savings through aggressive discharging [34], they often require bidirectional hardware and increase battery degradation. The proposed limited-V2G hybrid approach delivers competitive 45–62% reductions on a larger fleet (100 EVs) with greater practicality, outperforming traditional MILP methods [24] and matching deep learning-based systems [37].

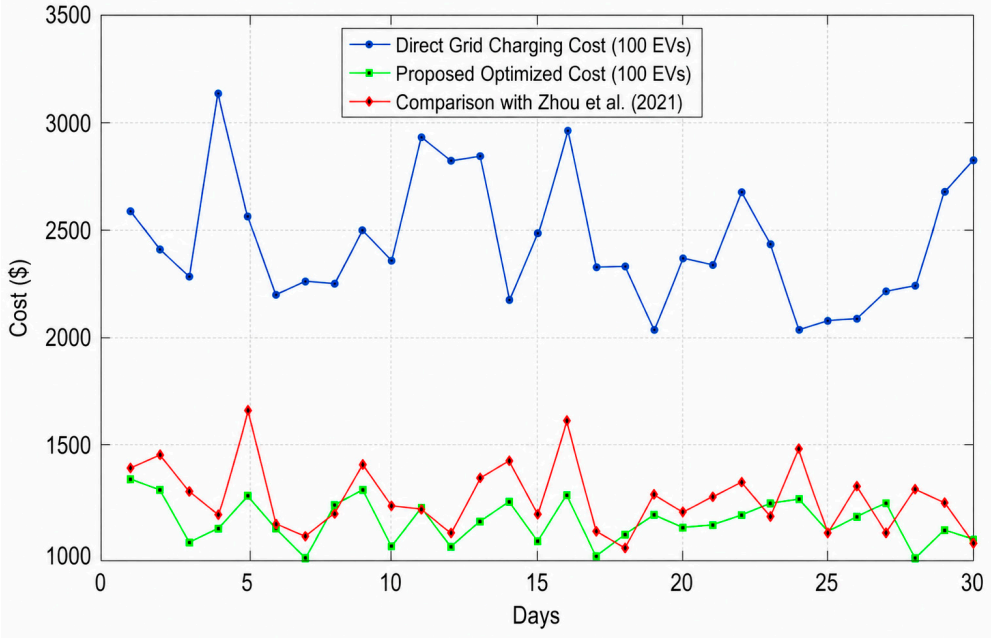

Figure 6. Monthly charging-cost comparison for a fleet of 100 EVs [24].

To ensure a fair baseline evaluation, the benchmark method in [24] was reproduced using the same dataset, time resolution, tariff signals, PV generation, and EV availability profiles adopted in this study. All preprocessing steps (chronological splitting, scaling/normalization, and 15 min discretization) were kept consistent to avoid confounding factors. In addition to the direct reproduced baseline comparison, a broader literature-level comparison is provided in Table 2 to contextualize the achieved cost reductions against reported results from other EV scheduling studies. It should be noted that reproducing multiple recent approaches (e.g., reinforcement learning and distributed optimization) on this dataset is not directly comparable without substantial redesign because those methods

typically depend on different system assumptions, objective definitions, uncertainty modeling, communication topology, and training environments. Therefore, this study prioritizes strict apples-to-apples benchmarking through a reproducible baseline implementation, complemented by literature comparison for broader context.

Table 2. Comparison of V2G-enabled EV scheduling approaches.

Study/Methodology	V2G Enabled	Baseline Used for Cost Saving	Reported Cost Savings	Key Remarks
Zhou et al. (2021) [24]	Yes	Uncoordinated charging	30–35%	Centralized MILP; commercial-building prosumers; deterministic optimization
Chopra et al. (2025) [35]	Yes	Uncoordinated charging	Up to 51.6%	Horse Herd Optimization; rooftop solar; techno-economic assessment
Onishi et al. (2020) [36]	Yes	Uncoordinated charging	Up to 42%	MILP; renewables (incl. solar) + TOU/DRP; microgrid focus
Guo et al. (2022) [38]	Yes	Uncontrolled charging	20–35%	Deep RL; solar parking lot; handles uncertainties
Qian et al. (2023) [39]	Yes	Uncoordinated charging	15–30% (peak load reduction + cost)	Multi-agent DRL; radial distribution network; focuses on load balancing + V2G profit
Ran et al. (2024) [37]	Yes	Uncoordinated charging	9.73% in total cost (plus peak-valley reduction)	Mixed-integer optimization; coupled power-transportation network; dynamic pricing; mobility-aware; peak shaving and stability focus
Lee et al. (2025) [40]	Yes	Uncoordinated charging	25.2–33.7%	Dynamic programming; TOU pricing; bidirectional focus; SOC-dependent
Shaheen et al. (2024) [41]	Yes	Uncoordinated charging	30–50%	Metaheuristics (incl. WOA best); TOU pricing; user cost minimization + grid support (peak/valley management); parking-based scheduling
Proposed PLEMS	Yes	Uncoordinated direct grid charging	45–62%	Hybrid MILP + GRU; scalable for n EVs; accurate forecasting; practical limited V2G

To provide a broader context, Table 2 compares the achieved cost savings with recent V2G-enabled works that use uncoordinated charging as the baseline. The proposed PLEMS achieves daily cost reductions of 45–62% compared to uncoordinated charging on a larger scale of 100 EVs. These savings are comparable to, or higher than, those reported in recent research relying on full V2G implementations (typically 20–52%). A key advantage of PLEMS is that discharge is triggered only under the bounded conditions of Equation (6) rather than the continuous, on-demand bidirectional operation required by full V2G. This restricts discharge to a small, predictable subset of high-SOC, idle, export-favorable intervals, which is what enables the reduced hardware and degradation burden discussed below. This design choice improves practicality by reducing battery degradation from frequent discharging cycles and lowering the need for expensive bidirectional hardware in charging infrastructure. As a result, the scheme offers robust economic benefits while remaining

more feasible for real-world deployment in existing parking facilities. Figure 6 illustrates consistent monthly cost benefits, confirming sustained performance over longer horizons. This robustness under varying daily conditions is a common advantage in optimized V2G scheduling frameworks [35,37], though the proposed hybrid method maintains higher average savings without relying on full bidirectional flows. Recent federated DRL methods (e.g., Qian et al. [39]) achieve ~83% reduction in power fluctuations (leading to indirect cost benefits of 15–25%), while the proposed PLEMS delivers higher direct cost savings (45–62%) with real-world validation and constraint guarantees. Similarly, Ran et al. (2024) [37] utilizes mixed-integer optimization in a coupled power–transportation network with dynamic pricing and mobility-aware prediction, achieving 9.73% direct cost savings plus 26% peak–valley reduction against uncoordinated charging. Lee et al. (2025) [40] apply dynamic programming for bidirectional V2G under TOU pricing with SOC-dependent constraints, delivering 25.2–33.7% average savings (up to 80.3% in favorable scenarios) and outperforming linear programming. Shaheen et al. (2024) [41] leverage metaheuristic algorithms (with WOA performing best) for parking-based V2G scheduling to minimize user daily costs and enable grid peak/valley management, yielding substantial reductions (30–50% range in comparable studies) over uncoordinated baselines. In comparison, the proposed PLEMS delivers higher direct cost savings (45–62%) with real-world validation, constraint guarantees, and scalability for larger EV fleets, highlighting its edge in practicality and economic impact.

Power utilization trends complement the cost analysis. Figure 7 shows daily power demand profiles under three scenarios: non-optimized, literature-based optimization, and the proposed objective function. The proposed approach consistently yields flatter demand curves with reduced peaks.

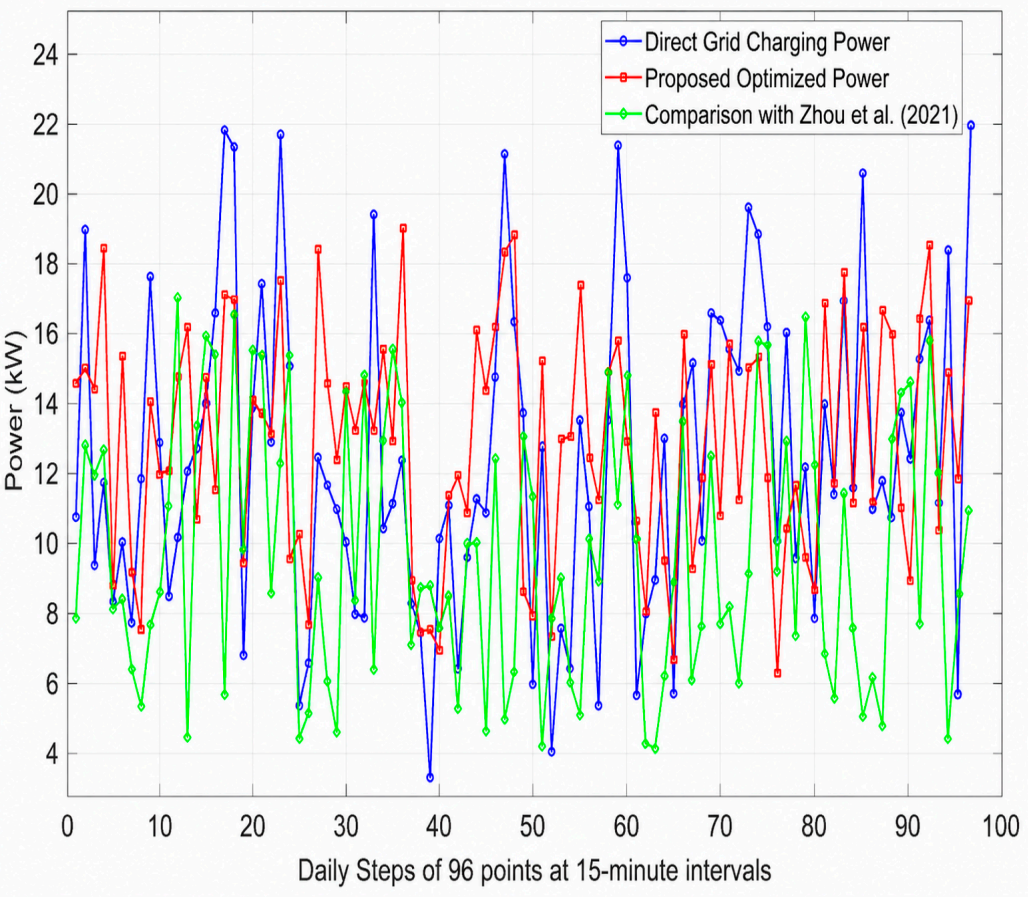

Figure 7. Power Utilization Trend of PLEMS for a Day [24].

Quantitative comparisons are summarized in Table 3, which reports average power, peak demand, and total energy consumption. The proposed method achieves the lowest average power (9.83 kW) compared to 13.16 kW (literature) and 13.53 kW (non-optimized), corresponding to reductions of 25.30% and 27.35%, respectively. Peak demand falls to 17.08 kW under the proposed approach, from 19.16 kW (literature) to 24.36 kW (non-optimized), representing 10.86% and 29.89% savings. Total energy consumption follows the same ranking, with the proposed method consuming 235.95 kWh versus 315.88 kWh (literature) and 324.81 kWh (non-optimized).

Table 3. Comparison of energy and power savings achieved by the proposed scheme.

Parameters	Proposed Scheme	Literature-Based Scheme [24]	Direct Grid Charging	Power Saving Compared to Literature-Based Based Scheme [24]	Power Saving Compared to the Direct Grid Charging
Average Power (kW)	9.83	13.16	13.53	25.32%	27.34%
Peak Power Consumption	17.08	19.16	24.36	10.86%	29.86%
Total Energy Consumption (kWh)	235.95	315.88	324.81	25.32%	27.35%

The scheduling phase is then validated using GRU-based SOC prediction. Figure 8 demonstrates close agreement between the GRU-predicted desired SOC and the MILP-derived optimal SOC values, confirming the accuracy of the optimization-guided SOC prediction approach. The main plot in Figure 8 shows the full-sequence predictions of the three architectures (RNN, GRU, and LSTM) compared against the ground-truth SOC values. Two zoomed insets demonstrate the superior long-term tracking capability of the GRU compared to the LSTM and RNN. The visual evidence aligns closely with the performance metrics in Table 3, verifying the superiority of the GRU model.

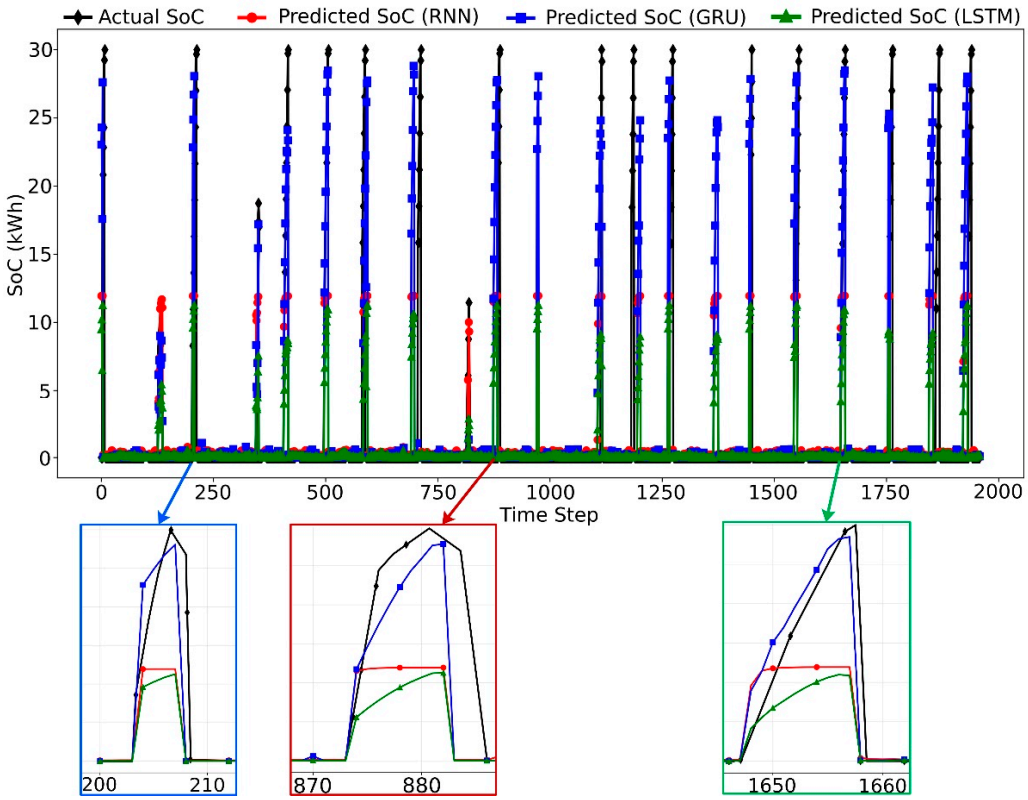

Figure 8. Predicted SOC using GRU in PLEMS.

Table 4 compares the performance of the proposed GRU-based model with LSTM and RNN on both training and test datasets. All metrics are computed using mean absolute error (MAE), root mean squared error (RMSE), and the coefficient of determination (R^2). The proposed GRU model achieves the lowest test RMSE of 0.9700 and the highest test R^2 of 0.9754, outperforming both the vanilla RNN (test RMSE = 1.1087, R^2 = 0.9678) and the LSTM (test RMSE = 1.7879, R^2 = 0.9163). Notably, the GRU also exhibits the smallest generalization gap between training and test performance, indicating superior robustness and better ability to capture the complex temporal dependencies in battery SOC data. The relatively poor performance of the LSTM suggests that its additional gating mechanism introduces unnecessary complexity for this particular dataset, leading to slight overfitting, whereas the GRU strikes an optimal balance between expressiveness and generalization. These results demonstrate that the proposed GRU architecture is particularly well-suited for accurate and reliable SOC prediction.

Table 4. Performance comparison of GRU, LSTM and RNN for SOC predictions.

Model	Train MAE	Train RMSE	Train R^2	Test MAE	Test RMSE	Test R^2
RNN	0.3657	1.2345	0.9600	0.3262	1.1087	0.9678
LSTM	0.4249	1.3524	0.9520	0.5204	1.7879	0.9163
GRU (Proposed)	0.2929	0.9795	0.9748	0.2938	0.9700	0.9754

Although predictive errors in recurrent models can potentially lead to error accumulation in closed-loop control, several aspects of PLEMS mitigate this risk. First, the GRU is trained on MILP-optimal trajectories that satisfy all physical constraints (SOC bounds, charge/discharge exclusivity, departure requirements), encouraging the model to produce feasible predictions. Second, in online operation, the system uses a step-by-step approach. At each 15 min interval, it calculates the charging or discharging power based on the latest measured SOC, instead of the previously predicted value. This naturally corrects any earlier prediction errors. Third, empirical evaluation shows very low prediction errors (test MAE of 0.29 kWh and RMSE of 0.97 kWh on 30 kWh batteries in Table 4), limiting accumulation over a daily horizon to negligible levels while maintaining constraint satisfaction and high cost savings. These factors collectively ensure stable and robust closed-loop performance under practical conditions.

Figure 9 illustrates charging and discharging patterns for selected EVs, where charging aligns with solar availability and low tariffs, and discharging events remain constrained to preserve user SOC requirements. The Figure 9 presents representative charging schedules and corresponding SOC profiles for four electric vehicles (EV7, EV38, EV71, and EV85) over a 24-hour period (96 time steps). The blue line represents the optimized charging/discharging power schedule, where positive values indicate charging and negative values indicate vehicle-to-grid (V2G) discharging. The orange line with circular markers shows the actual battery SOC obtained from the optimization results, while the purple dashed line with cross markers represents the SOC predicted by the proposed forecasting model. The close agreement between the orange and purple curves demonstrates the high prediction accuracy of the proposed model, ensuring that charging decisions closely follow the desired battery state while satisfying operational SOC constraints throughout the scheduling horizon.

The scheduling process aims to minimize grid stress and optimize energy consumption by distributing power efficiently among the EV fleet, as shown in Figure 10. The results demonstrate a clear reduction in power consumption for the scheduled EVs compared to

their non-scheduled counterparts. This highlights the efficacy of the scheduling approach in mitigating peak loads and optimizing energy utilization within the charging infrastructure.

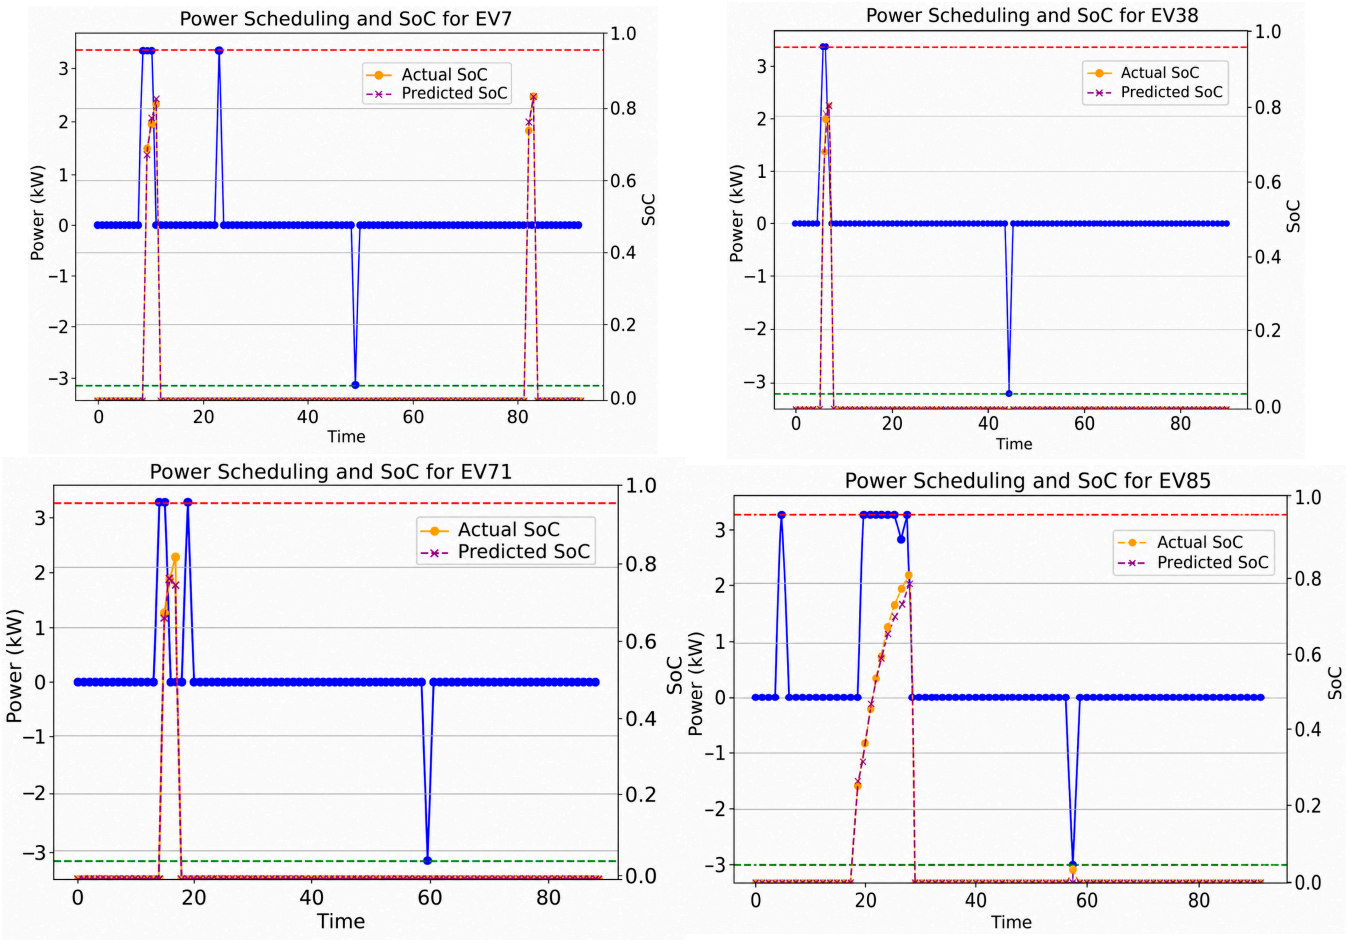

Figure 9. Charging and discharging of four randomly selected EVs through scheduling.

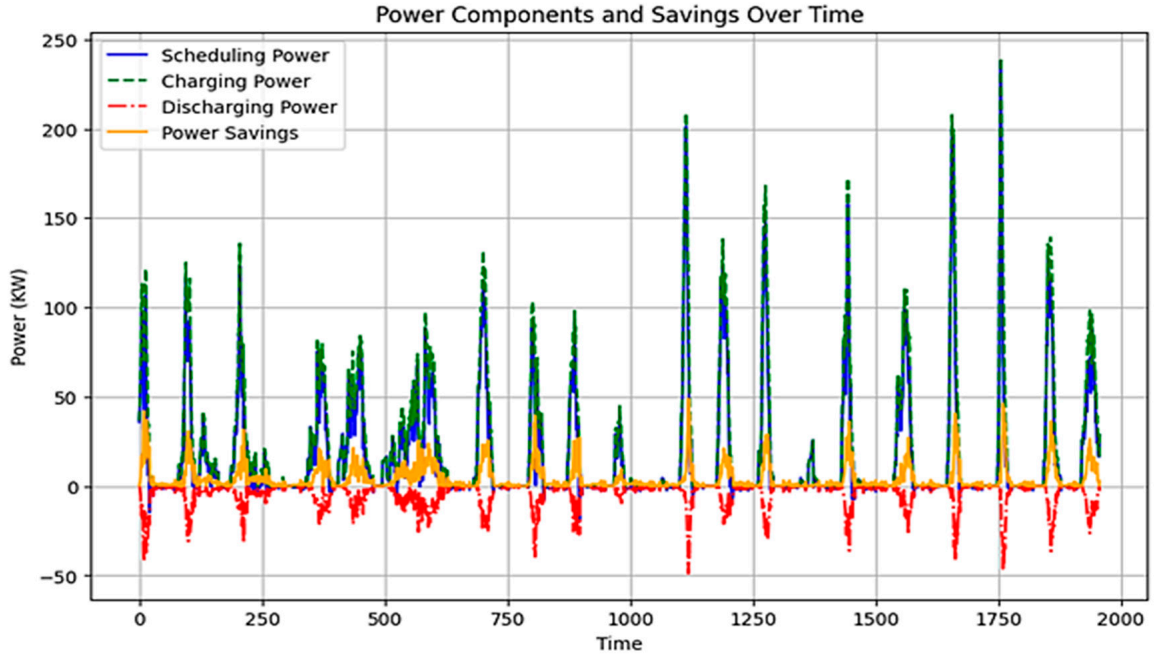

Figure 10. Comparison of power scheduling for 100 electric vehicles (EVs) over 30 days.

The scheduling bars for four selected EVs are represented in Figure 11, showcasing the charging energy utilization for each EV. Each red bar represents the charging energy without scheduling, and the blue bar shows charging energy with scheduling for an individual EV at different time intervals, giving a visual representation of the dynamic nature of charging demands. The fluctuations in power consumption are shown by varying heights of the bars for different EVs, reflecting the diverse charging requirements of the EV fleet. The purple broken (dashed) line in the Figure 11 represents the relative energy saving percentage achieved by the proposed scheduling for each selected EV. It connects the percentage labels (30.20% for EV7, 28.80% for EV38, 39.22% for EV71, and 61.66% for EV85) and serves as a visual trend line to highlight the varying degree of energy savings across the different vehicles. This line emphasizes how the scheduling is particularly effective for high-demand EVs (e.g., EV85 with the highest saving of 61.66%), while still providing substantial reductions for others. The dashed style distinguishes it from the solid bars representing actual energy consumption values.

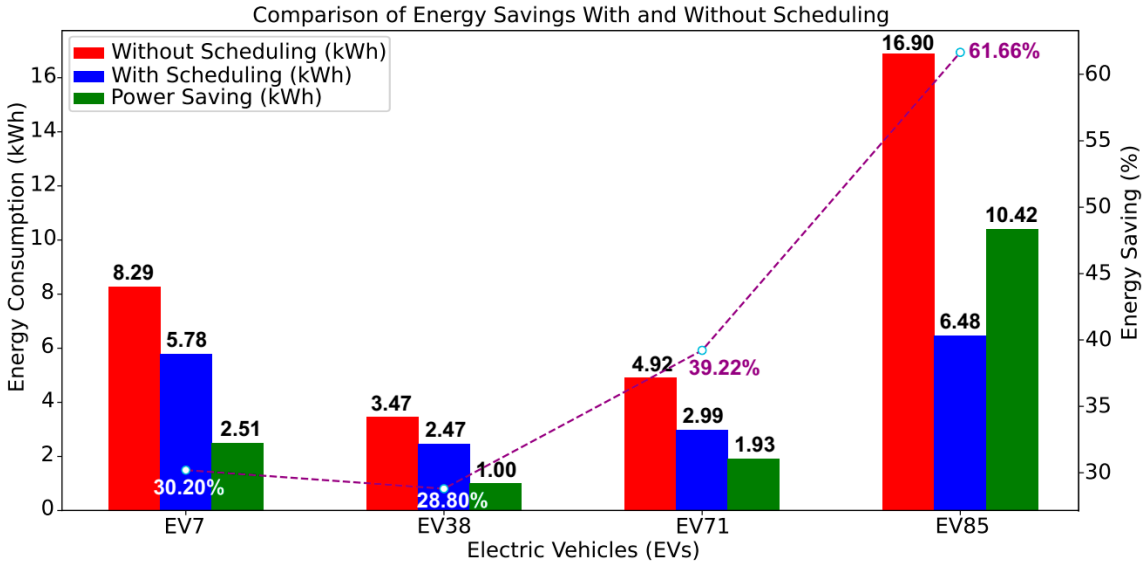

Figure 11. Power saving with proposed methodology.

A comparison of power scheduling for four EVs is provided in Table 5 in which each row represents an individual EV, detailing its scheduled power consumption (in kW) with and without scheduling. The difference in power consumption between the two scenarios is also given in Table 5. For instance, EV7 has a scheduled energy consumption of 5.78 kWh, whereas without scheduling, it consumes 8.28 kWh, resulting in an energy saving of 30%. Similarly, EV85 exhibits a notable reduction in energy consumption through scheduling, with an energy saving of 61.66%. Energy consumption for individual EVs is reduced by 28.8–61.66% under the proposed scheduling (Table 5). Such per-vehicle improvements align with findings in solar-integrated V2G systems, where intelligent coordination yields 30–50% energy expenditure reductions for high-demand vehicles [35,38].

Finally, statistical validation confirms that the improvements are not due to chance. Paired-sample *t*-tests were conducted over a full month of daily cost data (N = 30 days). Comparing daily non-optimized and proposed costs yields a t-statistic of 24.72 with a *p*-value below 0.0001, indicating strong significance. A second paired *t*-test between the proposed and literature-based [24] costs gives a t-statistic of 4.63 with a *p*-value of 0.00007, confirming additional gains. Both results are further supported by 95% confidence intervals, large effect sizes (Cohen's d), and non-parametric Wilcoxon signed-rank tests. The full summary of these results is presented in Table 6.

Table 5. Energy consumption and savings per EV.

EV ID	Energy Consumption (kWh) Without Scheduling	Energy Consumption (kWh) with Scheduling	Energy Saving (%)
EV7	8.28	5.78	30.2
EV38	3.47	2.47	28.8
EV71	4.92	2.99	39.3
EV85	16.9	6.48	61.66

Table 6. Statistical validation summary of proposed optimization approach (N = 30 days).

Metric	Comparison	N	Mean Saving	95% CI	Test Type	Test Statistic	*p*-Value	Effect Size (Cohen's d)	Result
Daily Cost	Non-optimized vs. Proposed	30	51.31%	48.98–53.64%	Paired *t*-test	t(29) = 24.72	<0.0001	4.51 (very large)	Statistically significant
Daily Cost	Non-optimized vs. Proposed	30	51.31%	48.98–53.64%	Wilcoxon signed-rank	W = 0.0	<0.0001	—	Confirmed
Daily Cost	Literature-based [24] vs. Proposed	30	6.76%	3.81–9.71%	Paired *t*-test	t(29) = 4.63	0.00007	0.85 (large)	Statistically significant
Daily Cost	Literature-based [24] vs. Proposed	30	6.76%	3.81–9.71%	Wilcoxon signed-rank	W = 52.0	0.00007	—	Confirmed

These statistical tests were conducted on N = 30 paired daily cost observations, corresponding to one full month of operation. Each paired observation compares the total daily charging cost under the proposed PLEMS scheduling against the corresponding day's cost under the comparison method (non-optimized or literature-based [24]), evaluated on the same day under matched demand and PV conditions.

For the non-optimized comparison, the mean daily saving was 51.31%, with a 95% confidence interval of 48.98–53.64% and a very large effect size (Cohen's d = 4.51). A paired *t*-test confirmed this difference was highly significant ($t(29) = 24.72$, $p < 0.0001$). A Shapiro–Wilk test supported the normality assumption underlying this test ($W = 0.952$, $p = 0.191$). A Wilcoxon signed-rank test, included as a distribution-free confirmation, corroborated the result ($W = 0.0$, $p < 0.0001$).

For the literature-based comparison, the mean daily saving was 6.76%, with a 95% confidence interval of 3.81–9.71% and a large effect size (Cohen's d = 0.85). A paired *t*-test confirmed this difference was statistically significant ($t(29) = 4.63$, $p = 0.00007$). A Shapiro–Wilk test likewise supported normality ($W = 0.978$, $p = 0.772$). A Wilcoxon signed-rank test corroborated the result ($W = 52.0$, $p = 0.00007$).

Together, these results confirm that the cost savings achieved by PLEMS over both baselines are statistically robust, supported by consistent evidence across a full month of daily operation, parametric and non-parametric testing, and adequate effect sizes.

Overall, the proposed two-phase framework (MILP optimization + GRU scheduling) consistently delivers economic and operational benefits:

- 51.5% daily cost reduction vs. non-optimized charging;
- 9.0% daily cost reduction vs. literature optimization [24];
- 27.3% lower average power consumption;
- 29.9% reduction in peak demand;
- 30–62% energy savings at the individual EV level.

These improvements support SDG7 (Affordable and Clean Energy) by reducing EV charging costs while increasing renewable energy integration. The results also demonstrate scalability and robustness, with the framework adapting to variations in tariffs and PV output.

While the proposed hybrid PLEMS framework demonstrates effective cost optimization and scalable real-time scheduling, several practical considerations should be noted. First, the learning-based scheduling stage relies on the representativeness and quality of historical data; significant changes in user behavior or operating conditions may require periodic retraining. Second, the performance of the system depends on the accuracy of PV generation and load forecasts, and forecast errors may affect optimal scheduling outcomes, although the closed-loop correction mechanism mitigates error propagation. Third, real-world deployment may involve additional considerations such as communication infrastructure, hardware integration, and implementation costs, which were not explicitly modeled in this study. Future work will investigate uncertainty-aware optimization, adaptive model updating, and deployment-oriented system integration to further enhance robustness and practical applicability.

In summary, while hybrid optimization-learning frameworks have been explored in energy systems, the proposed PLEMS advances their application to large-scale solar-integrated EV parking lot management through a practical two-stage design. The offline MILP generates constrained optimal labels for training a lightweight GRU, enabling accurate, decentralized real-time scheduling for 100 EVs without repeated optimization or central coordination. Validated on real historical Boulder data, this approach achieves high SOC forecasting accuracy (test R^2 = 0.975) and consistent daily cost savings of 45–62%, competitive with recent V2G methods while requiring only limited bidirectional capability for reduced implementation complexity and battery wear. These targeted improvements in scalability, computational efficiency, and deployment feasibility provide a balanced, practical contribution for sustainable urban EV charging infrastructure.

4. Conclusions and Future Work

This work proposed a PLEMS that optimizes the cost and scheduling of EV charging and discharging using a MILP framework. By integrating forecasted PV generation, predicted EV demand, and time-of-use pricing, the system minimizes charging costs while improving renewable energy utilization. The results demonstrated substantial cost savings (up to 51.86% compared to non-optimized charging), reduced peak demand, and enhanced grid stability. This cost reduction is the result of using on-site solar power, which the non-optimized baseline does not use at all, and the scheduling optimization performed by PLEMS. Separating these two effects is left for future work. Furthermore, the decentralized scheduling approach enables parking lot operators to manage energy flows autonomously, thereby supporting Sustainable Development Goal 7 (SDG7) on clean and affordable energy. While the proposed methodology effectively addresses PV variability and EV demand uncertainty through dynamic scheduling, challenges remain in transitioning from simulation to real-world deployment, including communication delays, infrastructure limitations, and the cost of advanced control and metering systems. Addressing these issues will require robust forecasting, flexible optimization, and adaptive control strategies to ensure reliable system operation under dynamic grid conditions. Future research should focus on real-world pilot testing in residential, commercial, and fleet contexts, robust and stochastic optimization to account for uncertainties in PV generation, EV arrivals, and electricity prices, as well as hybrid optimization approaches that combine MILP with metaheuristics or reinforcement learning for adaptive, data-driven decision-making. Additionally, integrating diverse renewable resources such as wind, bioenergy, and second-life EV batteries, along with techno-economic and policy analysis, will be crucial for enhancing resilience,

ensuring financial viability, and aligning with regulatory frameworks. Overall, the proposed PLEMS provides a scalable and cost-effective framework for sustainable EV charging, with its long-term success dependent on iterative real-world validation, interdisciplinary collaboration, and supportive policy measures that promote adoption and interoperability. While the current work validates the proposed framework on high-fidelity historical datasets, real-time Hardware-in-the-Loop tests and field pilots on commercial charging infrastructure are planned as the next phase of this research.

Author Contributions: Conceptualization, T.R. and A.R.B.; Methodology, T.R., A.R.B. and A.R.; Software, T.R., A.R.B. and M.F.; Validation, T.R., A.R.B., M.F., A.A. and A.R.; Formal Analysis, T.R., A.R.B., M.F., A.A. and A.R.; Investigation, T.R.; Resources, A.R.B. and M.F.; Writing—Original Draft Preparation, T.R.; Writing—Review and Editing, A.R.B., M.F., A.A. and A.R.; Visualization, T.R. and A.R.B.; Supervision, A.R.B. and A.R.; Project Administration, A.R.B. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding authors.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Saleem, S.; Farhan, M.; Raza, S.; Awan, F.G.; Butt, A.D.; Safdar, N. Power Factor Improvement and MPPT of the Grid-Connected Solar Photovoltaic System Using Nonlinear Integral Backstepping Controller. *Arab. J. Sci. Eng.* **2023**, *48*, 6453–6470.
2. Singh, A.R.; Vishnuram, P.; Alagarsamy, S.; Bajaj, M.; Blazek, V.; Damaj, I.; Rathore, R.S.; Al-Wesabi, F.N.; Othman, K.M. Electric vehicle charging technologies, infrastructure expansion, grid integration strategies, and their role in promoting sustainable e-mobility. *Alex. Eng. J.* **2024**, *105*, 300–330.
3. Sabyasachi, S.; Singh, A.R.; Godse, R.; Jaiswal, S.; Bajaj, M.; Srivastava, I.; Blazek, V.; Prokop, L.; Misak, S. Reimagining E-mobility: A holistic business model for the electric vehicle charging ecosystem. *Alex. Eng. J.* **2024**, *93*, 236–258. [CrossRef]
4. Pandey, K. Artificial Intelligence (AI) in Electric Vehicle Ecosystems: Challenges, Opportunities, and Models for Accelerated Adoption. *Int. J. Appl. Sci. Eng. Rev.* **2024**, *5*, 1–19. [CrossRef]
5. Qureshi, K.N.; Alhudhaif, A.; Jeon, G. Electric-vehicle energy management and charging scheduling system in sustainable cities and society. *Sustain. Cities Soc.* **2021**, *71*, 102990.
6. Zhang, K.; Chen, M.; Zhu, R.; Zhang, F.; Zhong, T.; Lin, J.; You, L.; Lü, G.; Yan, J. Integrating photovoltaic noise barriers and electric vehicle charging stations for sustainable city transportation. *Sustain. Cities Soc.* **2024**, *100*, 104996.
7. Al-Alwash, H.M.; Borcoci, E.; Vochin, M.-C.; Balapuwaduge, I.A.; Li, F.Y. Optimization schedule schemes for charging electric vehicles: Overview, challenges, and solutions. *IEEE Access* **2024**, *12*, 32801–32818. [CrossRef]
8. Alanazi, M.S. A MILP model for optimal renewable wind DG allocation in smart distribution systems considering voltage stability and line loss. *Alex. Eng. J.* **2022**, *61*, 5887–5901. [CrossRef]
9. Liu, L.; Wang, R.; Wang, Y.; Li, W.; Sun, J.; Guo, Y.; Qu, W.; Li, W.; Zhao, C. Comprehensive analysis and optimization of combined cooling heating and power system integrated with solar thermal energy and thermal energy storage. *Energy Convers. Manag.* **2023**, *275*, 116464.
10. Liu, P.; Liu, Z.; Fu, T.; Garg, S.; Kaddoum, G.; Hassan, M.M. Optimization of multi-vehicle charging and discharging efficiency under time constraints based on reinforcement learning. *Alex. Eng. J.* **2024**, *105*, 724–735. [CrossRef]
11. Appleton, S.; Fotouhi, A. A model-based battery charging optimization framework for proper trade-offs between time and degradation. *Automot. Innov.* **2023**, *6*, 204–219. [CrossRef]
12. Yin, W.; Ji, J. Research on EV charging load forecasting and orderly charging scheduling based on model fusion. *Energy* **2024**, *290*, 130126. [CrossRef]
13. Hong, J.; Liang, F.; Yang, H.; Zhang, C.; Zhang, X.; Zhang, H.; Wang, W.; Li, K.; Yang, J. Multi-forword-step state of charge prediction for real-world electric vehicles battery systems using a novel LSTM-GRU hybrid neural network. *eTransportation* **2024**, *20*, 100322. [CrossRef]
14. Sridharan, S.; Venkataraman, S.; Raja, S. SOC estimation in lithium ion batteries using a hybrid gated residual wavenet with LSTM-GRU network across varying temperatures. *J. Electrochem. Soc.* **2025**, *172*, 030510. [CrossRef]

15. Zhao, Y.; Li, Y.; Cao, Y.; Jiang, L.; Wan, J.; Rehtanz, C. An RNN with small sequence trained by multi-level optimization for SOC estimation in Li-ion battery applications. *IEEE Trans. Veh. Technol.* **2023**, *72*, 11469–11481. [CrossRef]
16. Gao, X.; Jia, K.; Gao, Z.; Xiao, S. Fractional-Order GRU Networks With Memory Units Based on Hausdorff Difference for SOC Estimations of Lithium-Ion Batteries. *IEEE Trans. Ind. Inform.* **2024**, *21*, 1576–1584.
17. Mousaei, A.; Naderi, Y.; Bayram, I.S. Advancing state of charge management in electric vehicles with machine learning: A technological review. *IEEE Access* **2024**, *12*, 43255–43283. [CrossRef]
18. Metwly, M.Y.; Ahmed, M.; Hamad, M.S.; Abdel-Khalik, A.S.; Hamdan, E.; Elmalhy, N.A. Power management optimization of electric vehicles for grid frequency regulation: Comparative study. *Alex. Eng. J.* **2023**, *65*, 749–760. [CrossRef]
19. Siddiqui, J.; Ahmed, U.; Amin, A.; Alharbi, T.; Alharbi, A.; Aziz, I.; Khan, A.R.; Mahmood, A. Electric Vehicle charging station load forecasting with an integrated DeepBoost approach. *Alex. Eng. J.* **2025**, *116*, 331–341. [CrossRef]
20. Xue, X.; Ai, X.; Fang, J.; Jiang, Y.; Cui, S.; Wang, J.; Ortmeyer, T.H.; Wen, J. Continuous approximate dynamic programming algorithm to promote multiple battery energy storage lifespan benefit in real-time scheduling. *IEEE Trans. Smart Grid* **2024**, *15*, 5744–5760. [CrossRef]
21. Liu, X.; Gou, Z. Hybrid forecasting and optimization framework for residential photovoltaic-battery systems: Integrating data-driven prediction with multi-strategy scenario analysis. In *Building Simulation*; Springer: Berlin/Heidelberg, Germany, 2025; pp. 1–23.
22. Dolgintseva, E.; Wu, H.; Petrosian, O.; Zhadan, A.; Allakhverdyan, A.; Martemyanov, A. Comparison of multi-step forecasting methods for renewable energy. *Energy Syst.* **2024**, *17*, 95–126. [CrossRef]
23. Wang, F.Y.; Chen, Z.; Hu, Z. Comprehensive optimization of electrical heavy-duty truck battery swap stations with a SOC-dependent charge scheduling method. *Energy* **2024**, *308*, 132773. [CrossRef]
24. Zhou, H.; Zhou, Y.; Hu, J.; Yang, G.; Xie, D.; Xue, Y.; Nordström, L. LSTM-based energy management for electric vehicle charging in commercial-building prosumers. *J. Mod. Power Syst. Clean Energy* **2021**, *9*, 1205–1216. [CrossRef]
25. Elghanam, E.; Abdelfatah, A.; Hassan, M.S.; Osman, A.H. Optimization techniques in electric vehicle charging scheduling, routing and spatio-temporal demand coordination: A systematic review. *IEEE Open J. Veh. Technol.* **2024**, *5*, 1294–1313. [CrossRef]
26. Zabihi, A.; Parhamfar, M. Decentralized energy solutions: The impact of smart grid-enabled EV charging stations. *Heliyon* **2025**, *11*, e41815. [CrossRef]
27. Wu, H.; Qi, Y.; Peng, J.; Tang, H.; Shao, M. STMO: A spatio-temporal multimodal optimization algorithm for optimizing electric vehicle charging scheduling and station layout. *Alex. Eng. J.* **2025**, *132*, 204–217. [CrossRef]
28. Sivanadipatham, C.S.; Singh, A.K.; Kumar, P.; Singh, S. A Comprehensive Decentralized Control Framework for Balanced and Reliable Residential EV Charging. *IEEE Trans. Transp. Electrif.* **2025**, *11*, 13014–13028. [CrossRef]
29. Sharma, S.; Ali, I. Efficient real-time cost management in renewable energy-powered microgrid with integrated electric vehicle charging/discharging control. *Majlesi J. Electr. Eng.* **2025**, *19*, 1–25.
30. Singh, A.R.; Kumar, R.S.; Rathore, R.S.; Pandian, A.; Alrayes, F.S.; Allafi, R.; Ahmad, N. AI-enhanced smart grid framework for intrusion detection and mitigation in EV charging stations. *Alex. Eng. J.* **2025**, *115*, 603–621. [CrossRef]
31. Singh, A.; Panigrahi, B.K. Prosumer Cost Efficiency and Ensuring Grid Stability through a Hierarchical PPO Framework in Decentralized Community Energy Management. *IEEE Access* **2025**, *13*, 38368–38386. [CrossRef]
32. Hasan, S.; Zeyad, M.; Ahmed, S.M.; Mahmud, D.M.; Anubhove, M.S.T.; Hossain, E. Techno-economic feasibility analysis of an electric vehicle charging station for an International Airport in Chattogram, Bangladesh. *Energy Convers. Manag.* **2023**, *293*, 117501. [CrossRef]
33. Rasheed, T.; Bhatti, A.R.; Farhan, M.; Rasool, A.; El-Fouly, T.H. Improving the efficiency of deep learning models using supervised approach for load forecasting of electric vehicles. *IEEE Access* **2023**, *11*, 91604–91619. [CrossRef]
34. El Harouri, K.; El Hani, S.; Naseri, N.; Aboudrar, I.; Daghouri, A. Optimizing electric vehicle charging in smart parking lots using particle swarm optimization: A comparative study in Morocco, France, and Tunisia. *IFAC J. Syst. Control* **2025**, *34*, 100338. [CrossRef]
35. Chopra, K.; Shah, M.; Niazi, K.; Sharma, G.; Bokoro, P.N. Optimizing Solar-Powered EV Charging: A Techno-Economic Assessment Using Horse Herd Optimization. *Energies* **2025**, *18*, 4556. [CrossRef]
36. Onishi, V.C.; Antunes, C.H.; Trovão, J.P.F. Optimal energy and reserve market management in renewable microgrid-PEVs parking lot systems: V2G, demand response and sustainability costs. *Energies* **2020**, *13*, 1884. [CrossRef]
37. Ran, Y.; Liao, H.; Liang, H.; Lu, L.; Zhong, J. Optimal Scheduling Strategies for EV Charging and Discharging in a Coupled Power–Transportation Network with V2G Scheduling and Dynamic Pricing. *Energies* **2024**, *17*, 6167. [CrossRef]
38. Guo, G.; Gong, Y. Energy management of intelligent solar parking lot with EV charging and FCEV refueling based on deep reinforcement learning. *Int. J. Electr. Power Energy Syst.* **2022**, *140*, 108061. [CrossRef]
39. Qian, J.; Jiang, Y.; Liu, X.; Wang, Q.; Wang, T.; Shi, Y.; Chen, W. Federated reinforcement learning for electric vehicles charging control on distribution networks. *IEEE Internet Things J.* **2023**, *11*, 5511–5525. [CrossRef]

40. Lee, H.; Kim, H.; Kim, H.; Kim, H. Optimal Vehicle-to-Grid Charge Scheduling for Electric Vehicles Based on Dynamic Programming. *Energies* **2025**, *18*, 1109. [CrossRef]
41. Shaheen, H.I.; Rashed, G.I.; Yang, B.; Yang, J. Optimal electric vehicle charging and discharging scheduling using metaheuristic algorithms: V2G approach for cost reduction and grid support. *J. Energy Storage* **2024**, *90*, 111816. [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.