Forecasting and Control of Peak Power at EV Charging Stations

Samuel Bobick
Narges Norouzi, Ed.

Electrical Engineering and Computer Sciences
University of California, Berkeley

Technical Report No. UCB/EECS-2025-135
http://www2.eecs.berkeley.edu/Pubs/TechRpts/2025/EECS-2025-135.html

June 3, 2025

Forecasting and Control of Peak Power at EV Charging Stations

By Samuel Bobick

Senior Honors Thesis
Department of Electrical Engineering and Computer Sciences
University of California, Berkeley
Spring 2025

Advisor: Professor Scott Moura
CS H195 Instructor of Record: Professor Narges Norouzi

Abstract

Workplace electric vehicle (EV) charging presents a promising opportunity to advance transportation electrification, reduce emissions, and improve charging equity. This thesis studies the workplace station-level joint price and power optimization problem, where the station operator optimizes a menu of price options to incentivize users to select controllable charging service.

In Chapter 1, we introduce the Smart Learning Pilot for EV Charging Stations (SLRP-EV), a cyber-physical testbed for joint price and power optimization on the UC Berkeley campus, and perform exploratory data analysis to convey usage patterns. In Chapter 2, we benchmark the performance of several forecasting approaches on the SLRP-EV dataset using metrics that evaluate model performance during peak-load events. Finally, in Chapter 3, we use the framework established in Chapters 1 and 2 to propose several solutions to enhance control of peak power. Through a Monte Carlo simulation, we find that model predictive control using a time series forecast significantly reduces peak power, demand charge, and overall operator costs.

This thesis is the culmination of 17 formative years of California public education, a profound gift bestowed upon me by society and brought to life by passionate, knowledgeable, and relatable teachers. It is with tremendous gratitude that I dedicate this thesis to the many educators who helped inspire my love for research, teaching, and learning.

With the strongest possible conviction, I call upon the leaders of this state, this country, and the world to safeguard and strengthen the institution of public education, a cornerstone of our democracy and a crucial foundation for the prosperity of future generations.

Contents

List of Figures

List of Tables

Chapter 1

Introduction to the Station-Level EV Charging Optimization Problem

1.1 Introduction

Smart electric vehicle (EV) charging at workplaces presents many key opportunities to contribute to transportation electrification and decarbonization. For example, although EV owners typically charge overnight, midday workplace charging capitalizes on cheap, renewable-laden electricity, thus realizing the full emissions reduction potential of EVs [23]. Also, workplace charging can alleviate inequities for EV owners without access to at-home chargers, such as those in multi-unit dwellings [21]. Furthermore, intelligent charging algorithms can help align charging with utility time-of-use (TOU) tariff schedules and shave peak power consumption to maximize profit and incentivize charger expansion by private companies [29]. EV charging stations are sometimes built to oversubscribe to grid resources. That is, the EV chargers would violate transformer constraints if they were all used at maximum power simultaneously. In this scenario load management is utilized to satisfy grid constraints [15]. Finally, EV charging station load management can be integrated with holistic demand flexibility that manage many types of loads at commercial sites [22].

This thesis studies the station-level joint price and power optimization problem. By incentivizing users with discounts if they allow their load to be controlled instead of uncontrolled, the station operator is able to increase revenue and reduce peak power, opening the door to the myriad of benefits described above. This work takes an intelligent systems approach, integrating tools from machine learning and control to improve joint price and power optimization at EV charging stations, with a particular focus on peak power reduction.

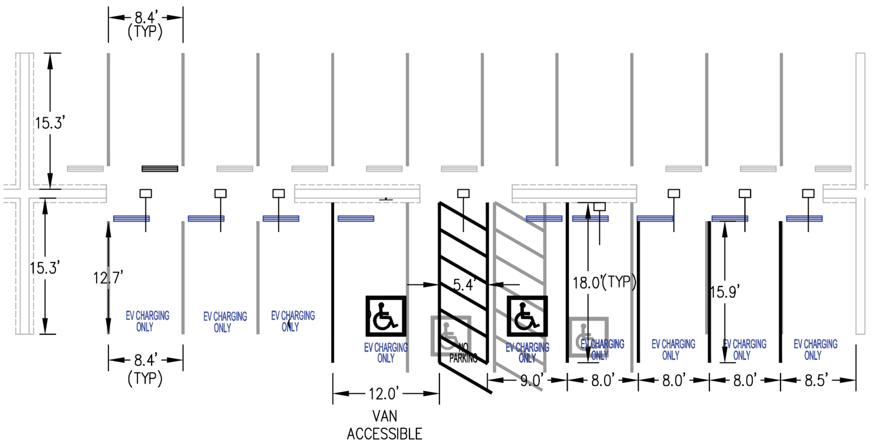

(a)

(b)

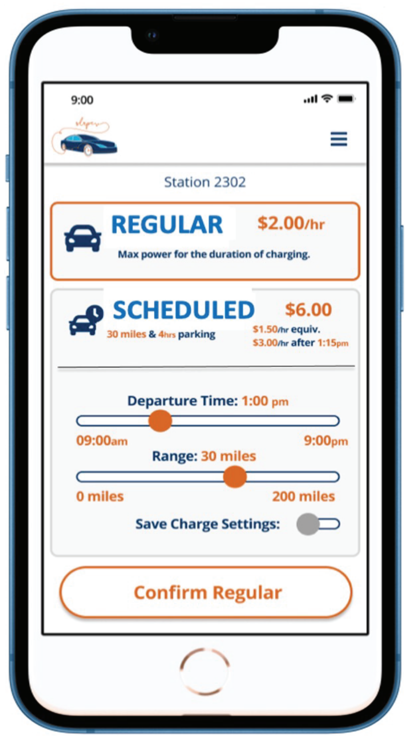

(c)

Figure 1.1: (a) UC Berkeley parking lot layout with EV spots in blue and non-EV spots in gray; (b) SLRP-EV chargers (6.6 kW); (c) Smartphone-based web app interface.

1.2 SLRP-EV: Smart Learning Pilot for EV Charging Stations

This study is conducted through the lens of the Smart Learning Pilot for EV Charging Stations (SLRP-EV), a workplace smart-charging facility located at UC Berkeley. The smart-charging framework attempts to maximize net revenue by jointly optimizing user-facing prices and power delivery.

Upon arrival at the charging station, the user inputs their required energy and expected time of departure from the charging station as illustrated in Figure 1.1(c). Given this information, the user is presented with two charging prices generated by a price and power optimization algorithm. Given these prices, the user can make one of three choices:

1. *REGULAR*: the battery charges at maximum power until it tops off or unplugs. Importantly, this load is uncontrollable.

2. *SCHEDULED*: the user's energy required is guaranteed to be provided by their indicated time of departure. The power delivery is scheduled with a control algorithm.

3. *LEAVE*: The user can leave if they consider the prices to be too high, opting for nearby chargers or a parking spot without a charger. Without this option, the price that maximizes station operator profit is infinite, which is non-sensical.

The only revenue source of SLRP-EV is the charging service fee collected from the users. The only cost is the electricity bill paid to the utility, which includes both TOU tariffs and demand charges.

The control objective is to maximize the expected profit from the station by choosing the *REGULAR* and *SCHEDULED* prices presented to the user, along with the power profiles for users who selected the *SCHEDULED* charging option.

1.3 Exploratory Data Analysis

In this section, we perform exploratory data analysis on the SLRP-EV dataset to convey usage patterns and forecasting challenges at the station level. The dataset contains information about the arrival and departure times of EVs at the station, energy requirements, prices presented, user choices, and power profiles at 15-minute increments.

Figure 1.2 shows the maximum station-level power across all hours of the week at SLRP-EV from 2021–2023. We observe much higher power peaks on weekdays. On weekdays, peak power is typically reached during the middle of the workday (i.e. 12–2 PM). Furthermore, peak power is highest during the UC Berkeley Fall and Spring semesters (i.e. January–May, August–December).

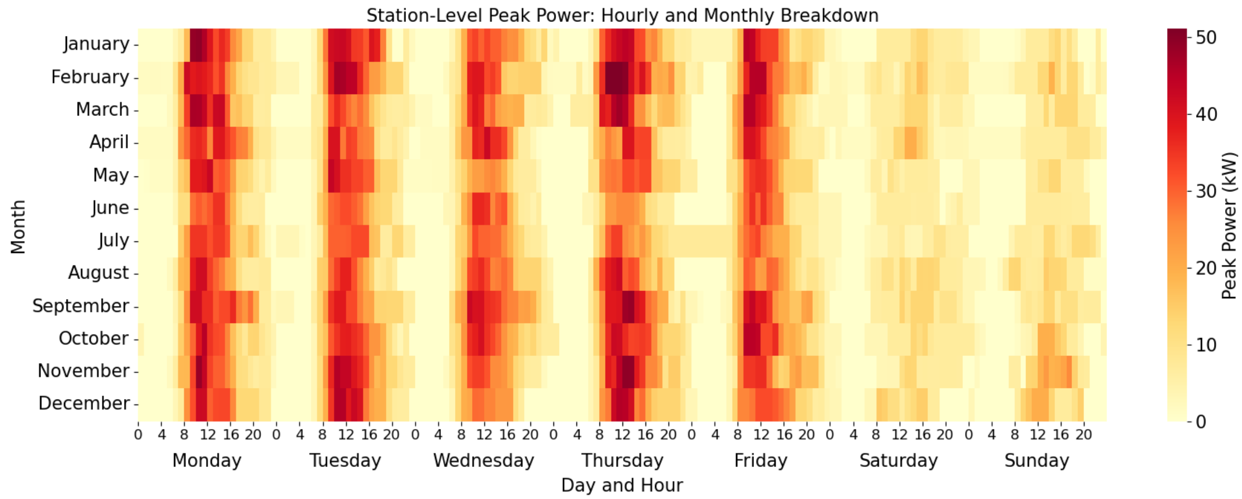

Figure 1.2: Heatmap showing peak power (kW) for each hour of the week at SLRP-EV from 2021–2023.

Figure 1.3 shows the relationship between peak power and key features from 2021–2023. A higher proportion of *SCHEDULED* sessions increases load flexibility and corresponds to lower peak power. Similarly, pricing *REGULAR* below *SCHEDULED* is linked to reduced peak power.

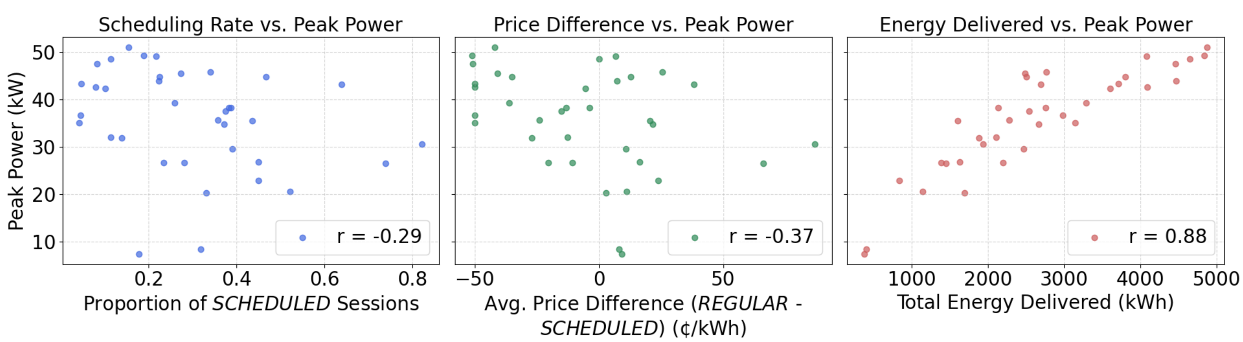

Figure 1.3: Scatter plots of peak power vs. several features. Each point represents 1 month (i.e. one billing cycle) from 2021–2023.

Figure 1.4 illustrates usage patterns at the SLRP-EV station. Arrivals peak between 6–10 AM, consistent with morning commute hours, while departures are concentrated in the afternoon. Charging activity is densest around midday (between arrival and departure peaks) and aligns with Pacific Gas & Electric's lowest TOU rate, "super off-peak," which is typically about half the cost of the "peak" rate [5].

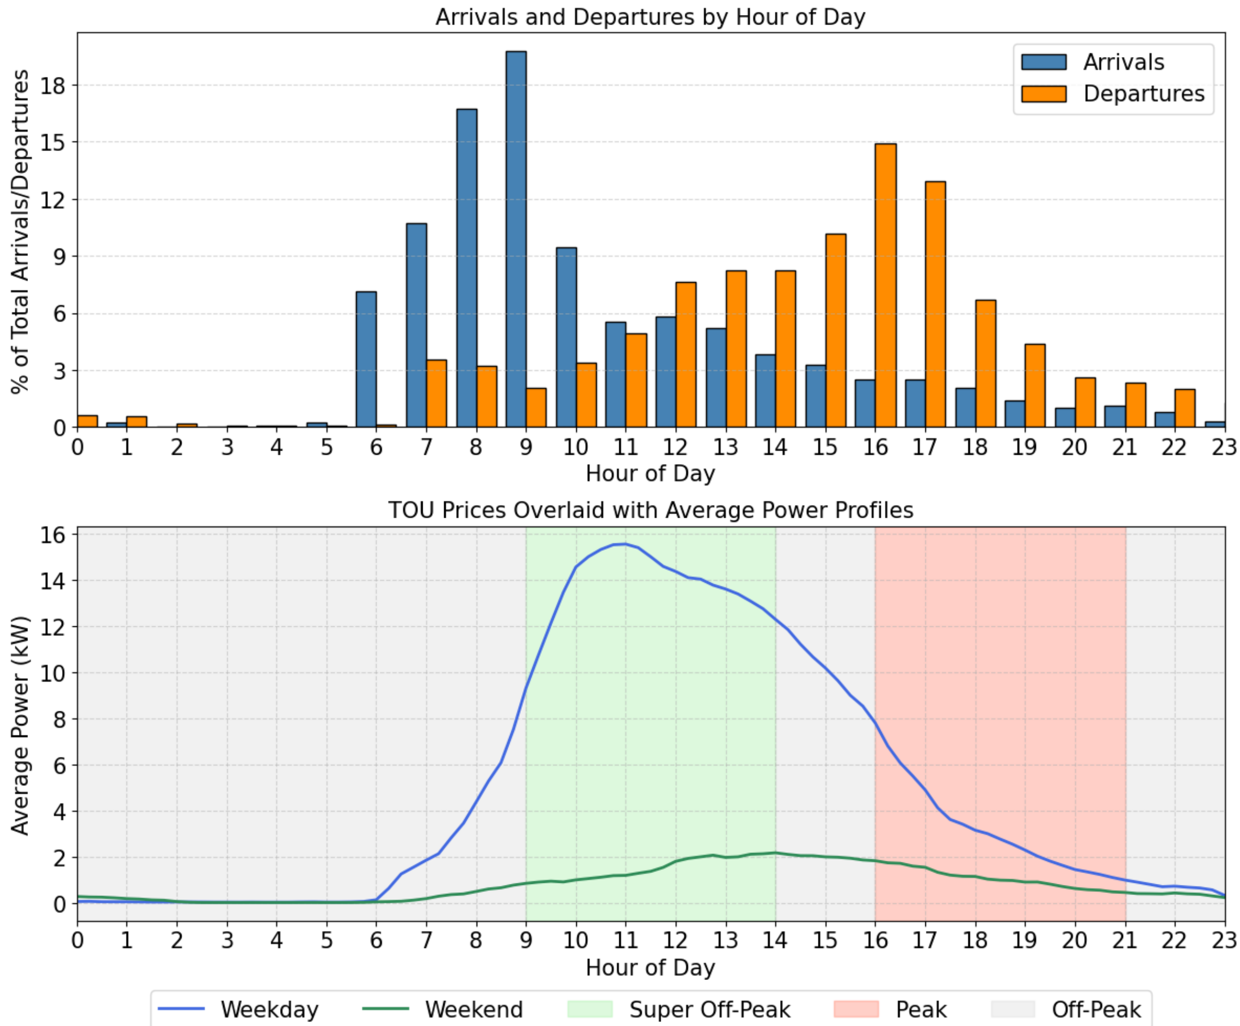

Figure 1.4: On the top panel, arrivals and departures broken down by hour of the day. On the bottom panel, average weekday and weekend power profiles overlaid with TOU tariff periods.

Chapter 2

Station-Level Load Forecasting

2.1 Introduction

Station-level load forecasting enables efficient, cost-effective, and grid-friendly operation of EV charging infrastructure. First, forecasts can be integrated into model predictive controllers, to be discussed in Chapter 3. Furthermore, forecasting enables coordination with other loads, avoiding overloads and aligning with grid capacity or grid signals (e.g. demand response) [25]. Finally, forecasting future load can enable EV charging stations to participate directly in electricity markets with vehicle-to-grid technology [26].

Jacob et. al. [10] and Nti et. al. [18] both provide a reviews of existing work on load forecasting. They differentiate between three main approaches: time series approaches [7, 1], classical machine learning approaches [8, 27], and deep learning approaches [33, 6, 12].

Station-level load forecasting at SLRP-EV involves just eight aggregated loads (i.e. eight chargers), significantly fewer than those considered in prior work [10, 18]. With such a small aggregation, the Law of Large Numbers has not yet taken effect. That is, the load curves are not as smooth and periodic as if there were more aggregated loads. This makes forecasting more challenging. In this chapter, we benchmark the performance of several forecasting approaches on the SLRP-EV dataset using metrics that reflect average, weighted, and peak prediction error.

2.2 Feature Engineering

Our models forecast the station load time series for the next l timesteps. Given the distinct load patterns on weekdays and non-weekdays (see Figure 1.4), we train separate models for each. Our models take the following as features:

1. The k most recent aggregate power values: $G_{T_{\text{start}}-k}, G_{T_{\text{start}}-k+1}, ...G_{T_{\text{start}}-1}$;

2. The planned station power profile for the next l timesteps for all active users in $\mathcal{I}$, considering that user n chooses $REGULAR$, and given power profiles $\bar{p}$ generated from

the previous optimization (triggered when user $n-1$ arrived) for users in $\mathcal{A}_{\text{sch}} \cup \mathcal{A}_{\text{reg}}$: $G_{T_{\text{start}}|REGULAR,\bar{p}}, \ldots G_{T_{\text{start}}+l|REGULAR,\bar{p}}$;

3. The number of active sessions at the charging station, including new user n: $|\mathcal{I}|$;
4. The total number of working chargers at the station;
5. A sinusoidal positional encoding of the time of day, time of week, and time of year [28];
6. A boolean variable that evaluates to 1 if the time index $T^{\text{start}} + 24$ hours falls on a workday.

2.3 Experiments

Models

We implement the following models[1]:

1. *Similar Day*: the average load at the same time and day over the past five weeks;
2. *Linear*: linear model trained with elastic net regularization [32];
3. *XGBoost* [4];
4. *FFNN*: a feed-forward neural network [2];

Loss Functions

We employ several loss functions. The first is standard RMSE. Given forecasted charging load $\hat{G}_\tau$ and ground truth G_τ,

$$RMSE = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(G_t - \hat{G}_t)^2}. \tag{2.1}$$

In the context to SLRP-EV, underprediction can lead to missed opportunities to reduce peak power via joint price and power optimization. As such, we adopt a weighted RMSE (WRMSE) loss to penalize underprediction. We use $\alpha = 2$.

$$\text{WRMSE} = \sqrt{\frac{1}{n}\sum_{\tau=1}^{n} w_\tau (G_\tau - \hat{G}_\tau)^2}, \quad w_\tau = \begin{cases} \alpha & \text{if } G_\tau > \hat{G}_\tau \\ 1 & \text{otherwise} \end{cases} \tag{2.2}$$

[1] Code available at `https://github.com/samuelBobick/StationLevelPowerForecasting`.

Since peak power prediction is of particular importance, both for satisfying grid constraints and managing demand charge, we also employ a weighted *peak* RMSE (WPRMSE), which is a weighted loss function which solely considers the peak power of each day, with $\alpha = 3$.

2.4 Results and Discussion

For each model, we perform a hyperparameter grid search, select the best configuration based on validation performance, then train and test the model, reporting test errors in Table 2.1. Each model uses $k = 96$ past 15-minute timesteps (24 hours of history) to predict $l = 32$ future timesteps (8 hours ahead). These parameters were selected for their effectiveness in the model predictive controller described in Chapter 3. Models are trained and tested on a random train-test split 2021–2023 SLRP-EV dataset.

Table 2.1: Error metrics for station-level load forecasting

Model	**RMSE (kW)**	**WRMSE (kW)**	**WPRMSE (kW)**
Similar Day	7.90	12.58	11.36
Linear	6.68	12.47	10.89
XGBoost	**6.50**	**12.13**	**10.57**
FFNN	6.96	12.83	11.22

Our results show that heuristic as simple as Similar Day yields forecasts that are reasonably accurate. However, models that employ the features detailed in Section 2.2 outperform Similar Day, not only in RMSE, but also in WRMSE and WPRMSE, which capture model performance during peak power events. Notably, a linear model using these features achieves competitive results with XGBoost while remaining interpretable. In the context of peak power control, a linear model is useful because it is convex in the features, which is useful for optimization. The marginal gains of XGBoost over the linear model are modest, suggesting that most of the predictive signal lies in well-engineered features rather than model complexity.

Chapter 3

Control Schemes for Peak Power Reduction

3.1 Introduction

A major cost for some EV charging station operators is demand charge, a utility fee based on the highest power consumption during a billing cycle. Approximately 5 million commercial customers in the United States are estimated to face retail electricity tariffs with demand charges greater than $15/kW, accounting for more than a quarter of the nation's 18 million commercial customers [17]. Demand charges are a significant component of commercial utility bills, generally accounting for 30% to 70% of total electricity costs [17]. An analysis in [9] finds that demand charge accounts for 70% to 94% of utility bills for DC fast charging stations in Southern California.

Demand charge reduction through smart charging presents a challenging control problem. Even with smart charging technology, EV charging station arrivals remain stochastic. On busy days, a single ill-timed arrival can push the station's load above the running peak, increasing demand charge.

Most prior work on demand charge mitigation for EV charging stations focuses on optimizing power schedules to maximize profit while accounting for EV arrival and departure times, energy demand constraints, demand charges, and TOU utility tariffs [14, 24, 11, 3, 30]. Lee, Pang, and Low [14] propose an offline pricing scheme that optimizes power delivery under fixed arrival, departure, and energy demand constraints. Their approach ensures cost recovery by incorporating facility demand charges, time-varying energy costs, and congestion costs. However, it does not account for stochastic variations in EV charging behavior. Yang et al. [30] use "block" model predictive control to optimize power delivery for demand charge mitigation under uncertain future arrivals, departures, and energy demands. However, their approach does not consider how to structure pricing to encourage user behavior

This chapter was co-authored with Thibaud Cambronne. T. Cambronne is with TotalEnergies S.E, Berkeley, CA 94704. Email: thibaud.cambronne@external.totalenergies.com.

that mitigates demand charge.

Formulations in [31, 19, 20] optimize both power delivery and user-facing prices. By incentivizing users with discounts if they allow their load to be controlled instead of uncontrolled, they are able to increase revenue. These formulations only control on-going charging sessions.

However, at workplace stations, most charging sessions begin in the morning and continue until the workday ends. Consequently, control decisions made at 8 AM can influence peak power at 2 PM. Early indicators of a peak-inducing day should be incorporated into morning control decisions. An effective controller needs anticipatory ability.

To solve this problem, we propose two solutions that provide anticipatory abilities: (i) convex reformulations of the demand charge constraint, and (ii) a forecast-enabled model predictive control (MPC) approach. Our contributions are as follows:

1. We propose an MPC method that anticipates demand spikes using time series forecasting.

2. We increase the load flexibility available to the controller by leveraging dynamic pricing to incentivize users to choose a controllable charging service.

3. We demonstrate a our approach outperforms the baseline algorithm and other non-MPC reformulations in reducing both demand charge and overall operator costs.

3.2 Baseline Optimization Problem

In this section, we present a baseline station-level optimization problem that is solved when a new user arrives at the charging station. Our formulation closely follows the formulation presented in [31], with some modifications.

Upon arrival at the charging station, the user inputs their required energy and expected time of departure from the charging station. Given this information, the user is presented with two charging prices generated by a price and power optimization algorithm. Given these prices, the user can make one of three choices:

1. *REGULAR*: the battery charges at maximum power until it tops off or unplugs. Importantly, this load is uncontrollable.

2. *SCHEDULED*: the user's energy required is guaranteed to be provided by their indicated time of departure. The power delivery is scheduled with a control algorithm.

3. *LEAVE*: The user can leave if they consider the prices to be too high, opting for nearby chargers or a parking spot without a charger. Without this option, the optimal service price that maximizes profit is infinite, which is non-sensical.

The control objective is to maximize the expected profit from the station by choosing the *REGULAR* and *SCHEDULED* prices presented to the user, along with the power profiles for users who selected the *SCHEDULED* charging option.

The baseline solution involves solving (3.1) directly. We assume that users choose one of the 3 charging options based solely on the price menu. We denote $\Pr(m \mid z)$ to be the probability that new user n chooses option m when presented with the set of prices z. We model these probabilities with a two-step discrete choice model, described in Appendix A. Due to the non-convexity of the discrete choice model (A.5) – (A.9), Eqn. (3.1) is difficult to solve numerically. As such, we find an approximate solution to (3.1) by searching over a grid of candidate prices $z_{\text{candidate}} = (z_{\text{sch}}, z_{\text{reg}})$. At each point in the grid, we solve $\min_p \mathbb{E}[f(p, m) | z_{\text{candidate}}]$, which is a linear program.

Objective Function

The joint price and power optimization problem is formulated as follows:

$$
\begin{aligned}
\min_{z,p}\ & \mathbb{E}[f(z, p, m)] && (3.1)\\
&= \Pr(m = \mathit{SCHEDULED} \mid z)\, f^{\text{sch}}(z, p^{\text{sch}}) && (3.2)\\
&+ \Pr(m = \mathit{REGULAR} \mid z)\, f^{\text{reg}}(z, p^{\text{reg}}) && (3.3)\\
&+ \Pr(m = \mathit{LEAVE} \mid z)\, f^{\text{leave}} && (3.4)\\
&\ \text{Demand charge constraints (3.5) – (3.7)},\\
&\ \text{Energy constraints (3.17) – (3.19)}.
\end{aligned}
$$

Suppose that a new user n arrives at time T^{start} and indicates a departure time of T_n^{end}. Then we optimize all *SCHEDULED* power profiles p over the optimization horizon $[T^{\text{start}}, T^{\text{end}})$, where $T^{\text{end}} = \max_{i \in \mathcal{I}} T_i^{\text{end}}$.

Peak Power Formulation

Peak power over a monthly billing cycle can be tracked with D_τ as follows, where the billing cycle starts at $\tau = 0$:

$$
\begin{aligned}
G_\tau &= \sum_{i \in A_{\text{sch}}} p_{i,\tau}^{\text{sch}} + \sum_{j \in A_{\text{reg}}} p^{\max} 1\{\tau < T_j^{\text{end}}\}, && (3.5)\\
D_0 &= 0, && (3.6)\\
D_{\tau+1} &= \max\{G_\tau, D_\tau\}. && (3.7)
\end{aligned}
$$

Scheduled Cost Function

Provided with new user n's energy required E_n^{req} and time of departure T_n^{end}, the cost function assuming that user n chooses the *SCHEDULED* charging option is given by

$$
\begin{aligned}
f^{\text{sch}}\left(z^{\text{sch}}, p^{\text{sch}}\right) = & \sum_{\tau=T^{\text{start}}}^{T_n^{\text{end}}-1} \left(c_\tau - z^{\text{sch}}\right) \cdot p_{n,\tau}^{\text{sch}} \cdot \Delta t && (3.8)\\
& + \sum_{i\in\mathcal{A}_{\text{sch}}} \left[\sum_{\tau=T^{\text{start}}}^{T_i^{\text{end}}-1} \left(c_\tau - \zeta_i\right) \cdot p_{i,\tau}^{\text{sch}} \cdot \Delta t\right] && (3.9)\\
& + \sum_{j\in\mathcal{A}_{\text{reg}}} \left[\sum_{\tau=T^{\text{start}}}^{T_j^{\text{end}}-1} \left(c_\tau - \zeta_j\right) \cdot p^{\max} \cdot \Delta t\right] && (3.10)\\
& + c_D \cdot \left(D_{T^{\text{end}}-1} - D_{T^{\text{start}}-1}\right). && (3.11)
\end{aligned}
$$

The key element of this equation for our study is (3.11), which represents the cost associated with the *increase in peak power* over the optimization horizon $[T^{\text{start}}, T^{\text{end}})$. Note that, as indicated by (3.7), this term cannot be negative and only becomes positive when the peak power increases over the optimization horizon.

Regular Cost Function

Provided the required energy E_n^{req} from the new user n, and the charge time T_n^{end} needed to fill the user's battery. The cost function assuming that user n chooses the *REGULAR* charging option is given by

$$
\begin{aligned}
f^{\text{reg}}\left(z^{\text{reg}}, p^{\text{reg}}\right) = & \sum_{\tau=T^{\text{start}}}^{T_n^{\text{end}}-1} \left(c_\tau - z^{\text{reg}}\right) \cdot p^{\max} \cdot \Delta t && (3.12)\\
& + \sum_{i\in\mathcal{A}_{\text{sch}}} \left[\sum_{\tau=T^{\text{start}}}^{T_i^{\text{end}}-1} \left(c_\tau - \zeta_i\right) \cdot p_{i,\tau}^{\text{sch}} \cdot \Delta t\right] && (3.13)\\
& + \sum_{j\in\mathcal{A}_{\text{reg}}} \left[\sum_{\tau=T^{\text{start}}}^{T_j^{\text{end}}-1} \left(c_\tau - \zeta_j\right) \cdot p^{\max} \cdot \Delta t\right] && (3.14)\\
& + c_D \cdot \left(D_{T_{\text{end}}-1} - D_{T_{\text{start}}-1}\right). && (3.15)
\end{aligned}
$$

Analogous to the scheduled case, (3.15) captures the cost associated with the increase in peak power over optimization horizon $[T^{\text{start}}, T^{\text{end}})$.

Leave Cost Function

If the user chooses to leave, the station operator incurs neither cost nor benefit. Thus,

$$f^{\text{leave}} = 0. \tag{3.16}$$

Energy Constraints

Constraints (3.17) – (3.19) ensure that the chosen power profiles satisfy the user's energy requirements.

$$\begin{aligned}
e_{i,\tau} &= e_{i,\tau-1} + \Delta t \cdot \eta \cdot p_\tau, && (3.17)\\
&\quad \text{for } \tau \in [T^{\text{start}}, T_i^{\text{end}}), \forall i \in \mathcal{I} \\
E_i^{\text{req}} &\leq e_{i,T^{\text{end}}}, \forall i \in \mathcal{I} && (3.18)\\
0 &\leq p_{i,\tau}^{\text{sch}} \leq p^{\max}, && (3.19)\\
&\quad \text{for } \tau \in [T^{\text{start}}, T_i^{\text{end}}), \forall i \in \mathcal{I}
\end{aligned}$$

3.3 Controllers for Peak Power Reduction

In this section, we propose several modifications to (3.1) which aim to reduce peak power. Section 3.3 applies a hard constraint to peak power. Section 3.3 gives extra weight to the demand charge term when the current power profile is on the verge of exceeding peak power. Section 3.3 uses time series forecasting to give the controller an anticipatory ability.

Iterative Hard Thresholding

In this approach, we add a hard constraint on peak power. That is, we solve (3.1) with the additional constraint

$$G_\tau \leq M \quad \text{for } \tau \in [T^{\text{start}}, T^{\text{end}}). \tag{3.20}$$

This constraint starts at the running peak in the billing cycle so far, $D_{T^{\text{start}}-1}$. If (3.1) with constraint (3.20) is infeasible, we iteratively loosen the constraint by increasing M by some step size $\epsilon > 0$, and re-solve until we find a solution.

Softplus Demand Charge Penalty

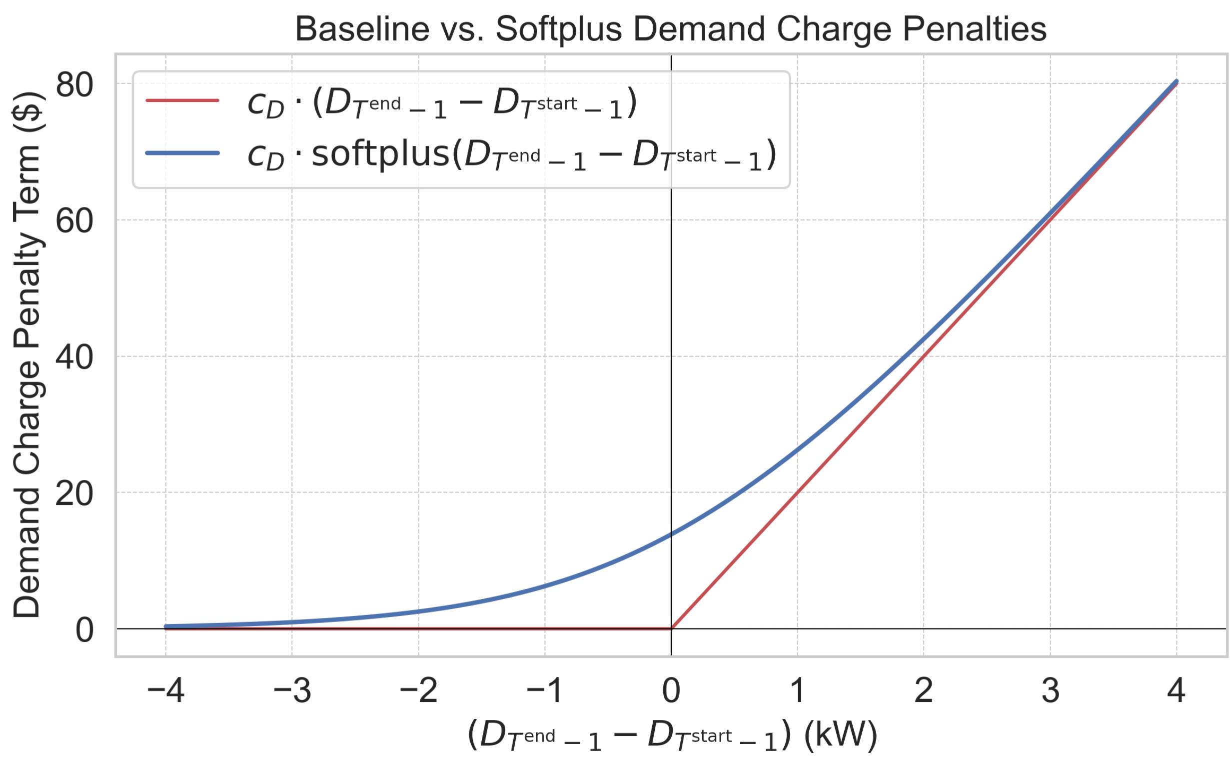

Figure 3.1: The softplus function adds a penalty when the peak of the optimized station power is approaching the running peak (as $x \to 0^-$), and converges to the baseline penalty as the station power exceeds the running peak (as $x \to +\infty$).

Approaching the running peak without surpassing it is problematic because any future arrival is likely to increase peak power. To add a demand charge penalty to a power profile that approaches the running peak, we apply the convex softplus function, $\text{softplus}(x) = \log(1+e^x)$, to the difference between $D_{T_{\text{start}}-1}$ and the peak power at the end of the optimization horizon. That is, we replace the terms in (3.11) and (3.15) with

$$c_D \cdot \text{softplus}\left(D_{T^{\text{end}}-1} - D_{T^{\text{start}}-1}\right). \tag{3.21}$$

Time Series Model Predictive Control

Next, we incorporate a time series forecast into the demand charge term to anticipate arrivals. Let $\{\hat{G}_{T_{\text{start}}|m}, \ldots, \hat{G}_{T_{\text{start}}+l|m}\}$ represent the forecasted station power time series for the next

l timesteps, given that new user n chooses charging option m. Then, we can replace terms (3.11) and (3.15) with

$$c_D \cdot \left(\max\{\hat{G}_{T^{\text{start}}|m}, ..., \hat{G}_{T^{\text{start}}+l|m}, D_{T^{\text{start}}-1}\} - D_{T^{\text{start}}-1} \right). \tag{3.22}$$

We test the MPC algorithm with 3 types of forecasters:

1. A naive model that assumes no additional arrivals, but uses, as a forecast, the sum of the optimal power profile from the *previous* optimization with user n's power profile, assuming they chose the $REGULAR$ option;
2. A linear model trained with elastic net regression [32];
3. XGBoost [4].

Integrating Forecasts into Timeseries MPC

Let ψ_τ represent the forecast of the station power profile, made outside of the control loop, using one of the methodologies presented in Chapter 2. At each solver iteration, to calculate $\hat{G}_{\tau|REGULAR}$ and $\hat{G}_{\tau|SCHEDULED}$ for timesteps $\tau \in [T^{\text{start}}, T^{\text{end}})$, we can simply start from the original forecast ψ_τ and add and subtract power profiles of the new control input p. Let $\bar{p}_i^{\text{sch}}$ represent the optimal scheduled power profile for user i found in the *previous* optimization triggered by user $n-1$. Then, we update our forecast as follows:

$$\hat{G}_{\tau|REGULAR} = \psi_\tau + \sum_{i \in \mathcal{A}_{\text{sch}}} \left(p_{i,\tau}^{\text{sch}} - \bar{p}_{i,\tau}^{\text{sch}} \right) \tag{3.23}$$

For each new user, the forecast is only executed once outside of the control loop, and is arithmetically updated to the new control actions at each iteration of the solver. Compared to executing the forecast inside the control loop, our single forecast approach allows us to:

1. Improve forecaster performance by using non-convex forecasters versus being limited to convex models, or having to increase computational time by using non-convex optimization;
2. Improve the forecast reliability by making predictions only on controlled historical power profiles versus candidate power profiles tested by the solver, which may not match the distribution of data that the forecaster was trained on;
3. Reduce computational time by making a single prediction instead of one per solver iteration.

Note that we post-process the forecast by clipping it to eliminate unrealistic forecasts:

$$\psi_\tau \geq p_{n,\tau}^{\text{reg}} + \sum_{i \in \mathcal{A}_{\text{sch}}} \bar{p}_{i,\tau}^{\text{sch}} + \sum_{i \in \mathcal{A}_{\text{reg}}} \bar{p}_{i,\tau}^{\text{reg}} \quad \forall \tau \in [T^{\text{start}}, T^{\text{end}}). \tag{3.24}$$

3.4 Numerical Study

Simulation Overview

We perform a Monte Carlo simulation[0] to quantitatively validate the performance of the control algorithms presented in the previous section. We replay all EV charging sessions sessions on the SLRP-EV platform from 2023, a total of 2274 charging sessions [20]. We assume a relatively high demand charge of $20/kW [17]. We use the Pacific Gas & Electric Commercial Business Electric Vehicle June 2023 rates for the TOU tariff structure [5]. Further details about the simulation setup can be found in Appendix B.

Table 3.1: Profit Comparison Across Control Schemes

Control Scheme	**Mean Cost/Revenue ($)**					**Change from Baseline (%)**		
	Demand Charge	**TOU**	**Revenue**	**Cost**	**Profit**	**Demand Charge**	**TOU**	**Cost**
Baseline	624	1,984	4,309	2,608	1,702	0.00	0.00	0.00
Threshold	628	1,987	4,331	2,615	1,716	-0.71	0.16	+0.29
Softplus	609	2,017	**4,446**	2,626	1,819	-2.34	1.67	+0.71
MPC (Naive)	**519**	1,970	4,253	**2,488**	1,765	**-16.90**	-0.71	**-4.58**
MPC (Linear)	521	1,970	4,362	2,491	1,870	-16.50	-0.68	-4.47
MPC (XGBoost)	533	**1,956**	4,434	2,489	**1,944**	-14.57	**-1.39**	-4.55

Results and Discussion

Table 3.1 presents the Monte Carlo simulation results, obtained by running each of the 12 months 10 times.

In particular, all MPC control schemes successfully decrease demand charge by approximately 15% from the baseline, while also decreasing TOU costs by approximately 1%, resulting in an approximately 4.5% average total cost reduction. The two non-MPC controllers fail to significantly reduce the total costs.

Figure 3.2 demonstrates how MPC improves peak power management. As the day's first sessions begin, the baseline controller waits to fulfill energy demand until the cheapest energy

[0] Simulation code available at `https://github.com/samuelBobick/StationLevelPowerForecasting`.

(super off-peak) is available. On the other hand, the MPC controller anticipates a busy day, and starts charging earlier so that there is less demand during the peak power event. In fact, by aggressively charging early in the day, the MPC algorithm is even able to partially avoid peak TOU pricing periods.

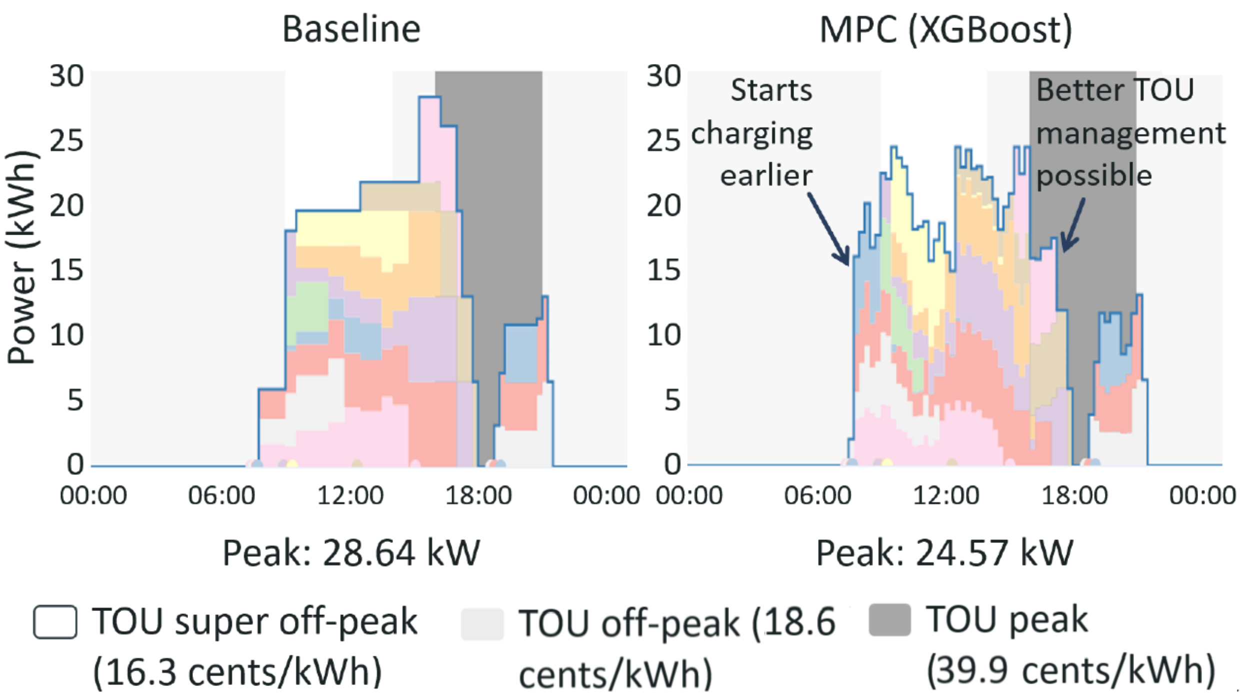

Figure 3.2: Example of control actions for baseline (left) and MPC (right) on September 6th 2023. On these two examples, the users all arrived at the same time, with the same requirements and all chose *SCHEDULED*. Each color represents a single user's power profile. The MPC algorithm shows better repartition of the load throughout the off peak hours and achieves a lower peak (24.57 kW) compared to the baseline solution (28.64 kW), while also reducing the energy consumed during peak TOU hours.

Figure 3.3 illustrates how the MPC controller leverages pricing to encourage users to select *SCHEDULED* when demand charge mitigation is pertinent. Both strategies raise prices before and during peak TOU hours when energy costs double. However, the MPC algorithm sets high prices early in the morning while significantly discounting *SCHEDULED* between 6-8 AM, prioritizing a base of controllable *SCHEDULED* sessions for the day. As a result, the MPC controller is able to maintain lower daytime prices, particularly for *REGULAR* sessions, whereas the baseline controller must raise prices during peak load hours to compensate for unanticipated sessions.

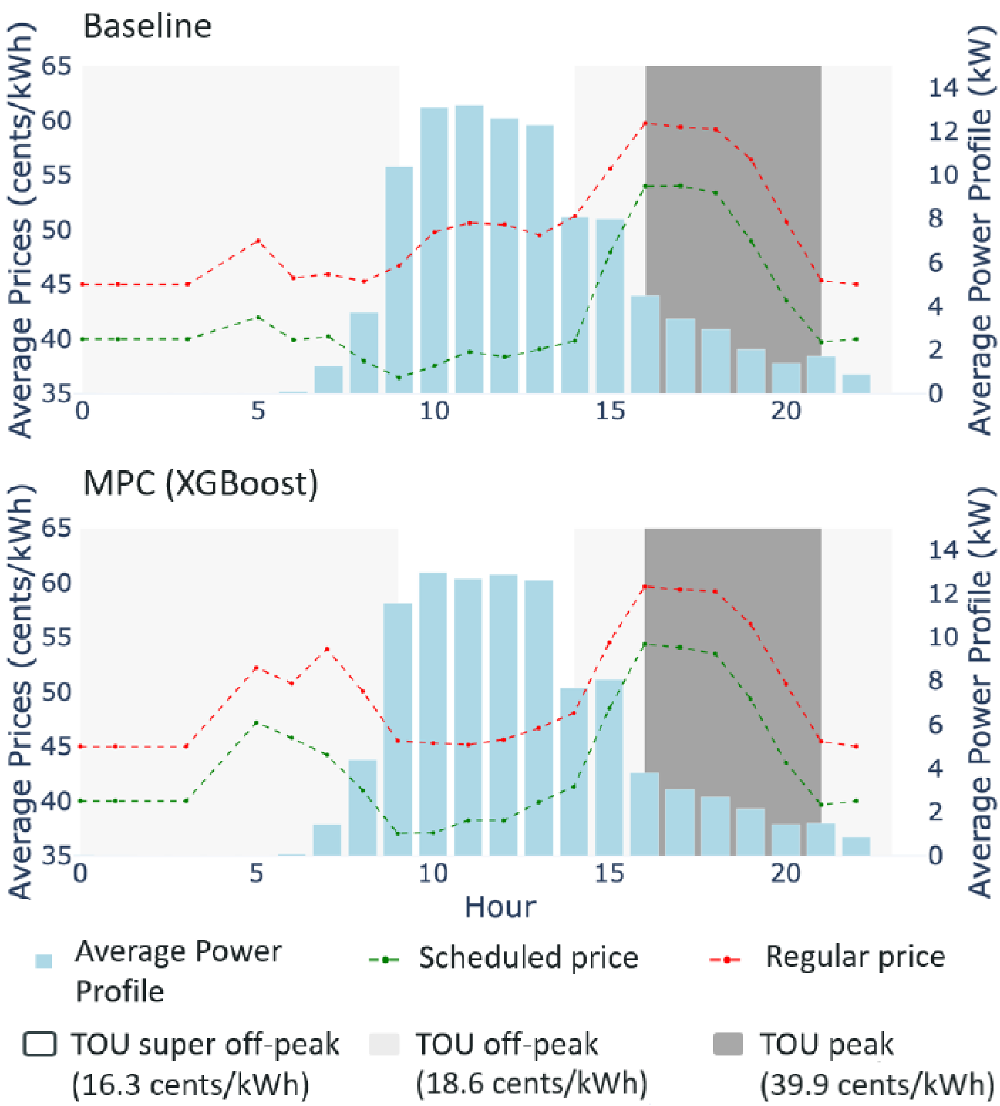

Figure 3.3: Example of pricing strategy for baseline (top) and MPC (bottom) solutions, overlaid with the average station power profile.

Forecast Performance vs. MPC Peak Power Reduction

Table 3.2 compares forecast errors to simulated mean peak power. The training RMSE closely correlates with the RMSE of predictions generated during the simulation. However, the naive forecast achieves the best peak power reduction despite a much higher training RMSE and a higher simulation RMSE. This suggests that minimizing RMSE alone does not strongly correspond to improved MPC performance.

Table 3.2: The Value of Forecast Accuracy in MPC

Control Scheme	**Training RMSE (kW)**	**Simulation RMSE (kW)**	**Mean Peak Power (kW)**
MPC (Naive)	11.36	7.46	**25.93**
MPC (Linear)	6.25	**4.13**	26.05
MPC (XGBoost)	**5.94**	4.31	26.65

One possible explanation for this is that the models were trained on data that does not exactly match the distribution of simulation data. Specifically, the training data contains charging sessions which were controlled with an optimization algorithm similar to the baseline algorithm. Retraining the model on data generated by the MPC controller during the simulation would improve the forecasts.

Our test case, the SLRP-EV station, has eight chargers. Due to its small size, station loads are stepwise and irregular, making forecasting more challenging. For a larger set of EV chargers, the Law of Large Numbers would smooth load curves, increasing periodicity and improving forecast accuracy. Consequently, the MPC approach may perform even better at larger stations or across a collection of aggregated station loads.

3.5 Conclusions and Limitations

In this paper, we present several methods to decrease demand charge in a joint price and power optimization scheme, and validate these experiments with Monte Carlo simulations. The MPC implementation detailed in Section 3.3 performs particularly well and decreases operator costs, especially demand charge.

In the baseline solution, the user whose arrival pushes the station load above its peak bears a disproportionate share of the demand charge costs, even though other users are already contributing to the peak. Our MPC approach addresses this issue by distributing the cost burden of demand charge expenses across all users whose usage is *forecasted* to contribute to an increase in peak power, ensuring a more fair allocation. A detailed analysis of how demand charge costs are allocated fairly across users via online pricing presents an potential avenue for future research.

In Section 3.4, we find that minimizing RMSE alone does not strongly correlate to improved MPC performance. Li, Ju, and Wang [16] have a similar finding in the context of building energy management. The authors find that forecasting errors have an asymmetric impact on MPC performance when demand charge is part of the objective. Further work is required to develop forecasting error metrics and techniques that integrate well with demand charge mitigation controllers.

When simulating SLRP-EV user sessions, we randomly select between *REGULAR* and *SCHEDULED* choices with probabilities specified by the discrete choice model. While we account for *LEAVE* in the optimization problem presented in (3.4), in the Monte Carlo simulation we are limited in the fact that users have already decided not to *LEAVE*. This limitation arises because the SLRP-EV dataset only contains information about users who chose to charge. Thus, profit and revenue should be interpreted with caution. Ultimately, station operators optimize for profit, so further work is needed to incorporate *LEAVE* behavior into the Monte Carlo simulation.

Ultimately, we find that MPC significantly decreases charging station operating costs, particularly demand charge. These savings can either (i) boost operator profits or (ii) allow the operator to offer lower prices. Enhancing behavioral simulation, testing our algorithms for control of large aggregate loads, and experimentally deploying our control schemes with real human behavior as in [20, 13] are key steps toward further improvement of joint price and power control for demand charge mitigation at workplace EV charging stations.

Bibliography

[1] Siddharth Arora and James W Taylor. "Short-term forecasting of anomalous load using rule-based triple seasonal methods". In: *IEEE transactions on Power Systems* 28.3 (2013), pp. 3235–3242.

[2] George Bebis and Michael Georgiopoulos. "Feed-forward neural networks". In: *Ieee Potentials* 13.4 (1994), pp. 27–31.

[3] Kalpesh Chaudhari et al. "Learning assisted demand charge mitigation for workplace electric vehicle charging". In: *IEEE Access* 10 (2022), pp. 48283–48291.

[4] Tianqi Chen and Carlos Guestrin. "Xgboost: A scalable tree boosting system". In: *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. 2016, pp. 785–794.

[5] *Electric Schedule Business Electric Vehicle*. Pacific Gas and Electric Company, 2023. URL: `https://www.pge.com/assets/rates/tariffs/CommlElecVehicle_230601-230630.xlsx`.

[6] Krzysztof Gajowniczek and Tomasz Zabkowski. "Short term electricity forecasting using individual smart meter data". In: *Procedia Computer Science* 35 (2014), pp. 589–597.

[7] Stephen Haben et al. "Short term load forecasting and the effect of temperature at the low voltage level". In: *International Journal of Forecasting* 35.4 (2019), pp. 1469–1484.

[8] Samuel Humeau et al. "Electricity load forecasting for residential customers: Exploiting aggregation and correlation between households". In: *2013 Sustainable internet and ICT for sustainability (SustainIT)*. IEEE. 2013, pp. 1–6.

[9] Rocky Mountain Institute. *EVgo Fleet and Tariff Analysis*. Tech. rep. Rocky Mountain Institute, 2017. URL: `https://rmi.org/wp-content/uploads/2017/04/eLab_EVgo_Fleet_and_Tariff_Analysis_2017.pdf`.

[10] Maria Jacob, Cláudia Neves, and Danica Vukadinović Greetham. *Forecasting and assessing risk of individual electricity peaks*. Springer Nature, 2020.

[11] Emre Can Kara et al. "Estimating the Benefits of Electric Vehicle Smart Charging at Non-Residential Locations: A Data-Driven Approach". In: *CoRR* abs/1503.01052 (2015). URL: `http://arxiv.org/abs/1503.01052`.

[12] Weicong Kong et al. "Short-term residential load forecasting based on LSTM recurrent neural network". In: *IEEE transactions on smart grid* 10.1 (2017), pp. 841–851.

[13] George Lee et al. "Adaptive charging network for electric vehicles". In: IEEE, Dec. 2016, pp. 891–895. ISBN: 978-1-5090-4545-7. DOI: 10.1109/GlobalSIP.2016.7905971.

[14] Zachary J Lee, John ZF Pang, and Steven H Low. "Pricing EV charging service with demand charge". In: *Electric Power Systems Research* 189 (2020), p. 106694.

[15] Zachary J Lee et al. "Adaptive charging networks: A framework for smart electric vehicle charging". In: *IEEE Transactions on Smart Grid* 12.5 (2021), pp. 4339–4350.

[16] Lunlong Li, Yi Ju, and Zhe Wang. "Quantifying the impact of building load forecasts on optimizing energy storage systems". In: *Energy and Buildings* 307 (2024), p. 113913.

[17] Joyce A McLaren, Pieter J Gagnon, and Seth Mullendore. *Identifying potential markets for behind-the-meter battery energy storage: A survey of US demand charges*. Tech. rep. National Renewable Energy Lab.(NREL), Golden, CO (United States), 2017.

[18] Isaac Kofi Nti et al. "Electricity load forecasting: a systematic review". In: *Journal of Electrical Systems and Information Technology* 7 (2020), pp. 1–19.

[19] Hassan Obeid et al. "Learning and Optimizing Charging Behavior at PEV Charging Stations: Randomized Pricing Experiments, and Joint Power and Price Optimization". In: *Applied Energy* 351 (2023), p. 121862. DOI: 10.1016/j.apenergy.2023.121862.

[20] Ayşe Tuğba Öztürk et al. "Joint Price and Power Optimization Experiment for Workplace Charging Stations". Under review at Sustainable Cities and Society. 2025.

[21] David Peterson. "Addressing challenges to electric vehicle charging in multifamily residential buildings". PhD thesis. 2011.

[22] Shanti Pless et al. *Integrating electric vehicle charging infrastructure into commercial buildings and mixed-use communities: Design, modeling, and control optimization opportunities*. Tech. rep. National Renewable Energy Lab.(NREL), Golden, CO (United States), 2020.

[23] Siobhan Powell et al. "Charging infrastructure access and operation to reduce the grid impacts of deep electric vehicle adoption". In: *Nature Energy* 7.10 (2022), pp. 932–945.

[24] Siobhan Powell et al. "Controlled workplace charging of electric vehicles: The impact of rate schedules on transformer aging". In: *Applied Energy* 276 (Oct. 2020), p. 115352. ISSN: 03062619. DOI: 10.1016/j.apenergy.2020.115352.

[25] Miadreza Shafie-Khah et al. "Optimal behavior of electric vehicle parking lots as demand response aggregation agents". In: *IEEE Transactions on Smart Grid* 7.6 (2015), pp. 2654–2665.

[26] Harun Turker and Seddik Bacha. "Optimal Minimization of Plug-In Electric Vehicle Charging Cost With Vehicle-to-Home and Vehicle-to-Grid Concepts". In: *IEEE Transactions on Vehicular Technology* 67.11 (2018), pp. 10281–10292. DOI: 10.1109/TVT.2018.2867428.

[27] Oleg Valgaev, Friederich Kupzog, and Hartmut Schmeck. "Designing k-nearest neighbors model for low voltage load forecasting". In: *2017 IEEE Power & Energy Society General Meeting*. IEEE. 2017, pp. 1–5.

[28] A Vaswani. "Attention is all you need". In: *Advances in Neural Information Processing Systems* (2017).

[29] Brett Williams and JR DeShazo. "Pricing workplace charging: financial viability and fueling costs". In: *Transportation Research Record* 2454.1 (2014), pp. 68–75.

[30] Lei Yang et al. "EV Charging Scheduling Under Demand Charge: A Block Model Predictive Control Approach". In: *IEEE Transactions on Automation Science and Engineering* 21.2 (2024), pp. 2125–2138. DOI: `10.1109/TASE.2023.3260804`.

[31] Teng Zeng et al. "Inducing Human Behavior to Maximize Operation Performance at PEV Charging Station". In: *IEEE Transactions on Smart Grid* 12.4 (2021), pp. 3353–3363. DOI: `10.1109/TSG.2021.3066998`.

[32] Hui Zou and Trevor Hastie. "Regularization and variable selection via the elastic net". In: *Journal of the Royal Statistical Society Series B: Statistical Methodology* 67.2 (2005), pp. 301–320.

[33] Thierry Zufferey et al. "Forecasting of smart meter time series based on neural networks". In: *Data Analytics for Renewable Energy Integration: 4th ECML PKDD Workshop, DARE 2016, Riva del Garda, Italy, September 23, 2016, Revised Selected Papers 4*. Springer. 2017, pp. 10–21.

Appendix A

Behavioral Model Formulation

Our behavioral model describes the probability that user n selects charging choice m when presented with prices z. Since z represents prices per unit of energy ($/kWh), we multiply it by the maximum charging power $p^{\max}$ to convert it to a price per unit of time ($/hr). This transformation ensures consistency between the optimization framework in Section 3.2 and the utility model in [20].

$$z'_{\text{reg}} = z^{\text{reg}} p^{\max}, \quad z'_{\text{sch}} = z^{\text{sch}} p^{\max} \tag{A.1}$$

To estimate the choice probabilities, we use a two-step discrete choice model, using utilities U estimated in [20].

$$U_{\text{reg}} = 0.341 - 0.0184 \frac{(z'_{\text{reg}} - z'_{\text{sch}})}{2} \tag{A.2}$$

$$U_{\text{sch}} = 0.0184 \frac{(z'_{\text{reg}} - z'_{\text{sch}})}{2} \tag{A.3}$$

$$U_{\text{leave}} = -1 + 0.005 \frac{(z'_{\text{reg}} + z'_{\text{sch}})}{2} \tag{A.4}$$

First, we calculate the probability that the user chooses the *LEAVE* option:

$$\Pr(m = \mathit{LEAVE} \mid z) = \frac{e^{U_{\text{leave}}}}{e^{U_{\text{sch}}} + e^{U_{\text{reg}}} + e^{U_{\text{leave}}}}. \tag{A.5}$$

Then, we calculate the probability that the user chooses *SCHEDULED*, given that they do not choose the *LEAVE* option:

$$\Pr(m = \mathit{SCHEDULED} \mid z, m \neq \mathit{LEAVE}) \tag{A.6}$$

$$= \frac{e^{U_{\text{sch}}}}{e^{U_{\text{sch}}} + e^{U_{\text{reg}}}} \cdot (1 - \Pr(m = \mathit{LEAVE} \mid z)). \tag{A.7}$$

Likewise, for *REGULAR*,

$$\Pr(m = REGULAR \mid z, m \neq LEAVE) \tag{A.8}$$

$$= \frac{e^{U_{\mathrm{reg}}}}{e^{U_{\mathrm{sch}}} + e^{U_{\mathrm{reg}}}} \cdot (1 - \Pr(m = LEAVE \mid z)). \tag{A.9}$$

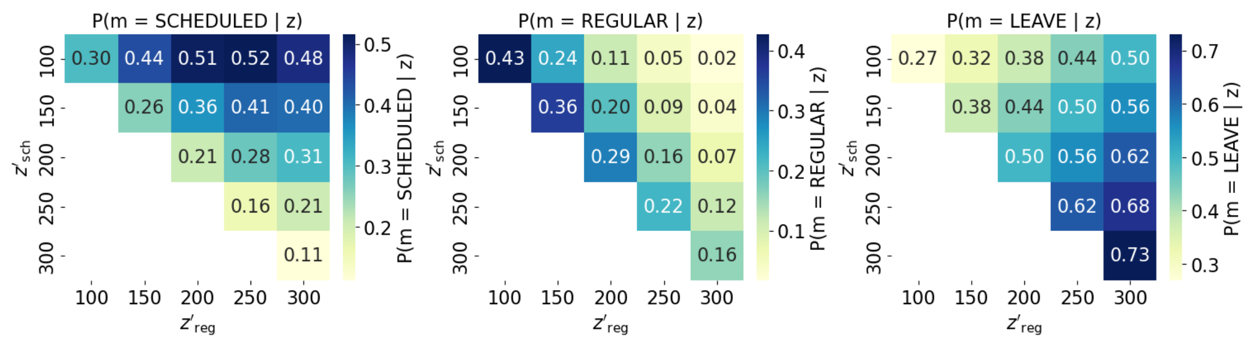

Figure A.1: Discrete choice probability distribution over a grid of prices.

Appendix B

Simulation Setup

The Slrp-EV dataset used to derive our simulation has a mix of *REGULAR* and *SCHEDULED* charging sessions. As we re-generate user choices using the discrete choice model, we make several assumptions:

1. Users have already made the decision to charge. That is, they cannot *LEAVE*. They only choose between *REGULAR* and *SCHEDULED* options.

2. The energy demand for charging sessions that were in reality *REGULAR* but are randomly simulated to be *SCHEDULED* have an energy demand equal to 57% of the energy delivered in the original charging session. We choose 57% because in the Slrp-EV dataset on average the E^{req} for *SCHEDULED* sessions is 57% of the energy consumed from a *REGULAR* session of the same length.

3. We assume we know that user n's time of departure is fixed and known.

4. We assume that all EVs can charge at the charger's maximum power rating ($p^{\max}$) of 6.6 kW.

Additional parameter settings are listed in Table B.1.

Table B.1: Simulation Parameters

Parameter	**Value**	**Description**
Δt	0.25 hours	Length of each time step
$p^{\max}$	6.6 kW	Maximum charging power
η	1	Charger efficiency
ϵ	1 kW	Threshold increment for solution 3.3
k	96	Number of past time steps to use as features for the timeseries forecaster in the MPC solution
l	32	Number of time steps the timeseries forecaster in the MPC solution