

Learning and Optimization for Price-based Demand Response of Electric Vehicle Charging

Chengyang Gu¹, Yuxin Pan², Ruohong Liu¹ and Yize Chen¹

Abstract—In the context of charging electric vehicles (EVs), the price-based demand response (PBDR) is becoming increasingly significant for charging load management. Such response usually encourages cost-sensitive customers to adjust their energy demand in response to changes in price for financial incentives. Thus, to model and optimize EV charging, it is important for charging station operator to model the PBDR patterns of EV customers by precisely predicting charging demands given price signals. Then the operator refers to these demands to optimize charging station power allocation policy. The standard pipeline involves offline fitting of a PBDR function based on historical EV charging records, followed by applying estimated EV demands in downstream charging station operation optimization. In this work, we propose a new decision-focused end-to-end framework for PBDR modeling that combines prediction errors and downstream optimization cost errors in the model learning stage. We evaluate the effectiveness of our method on a simulation of charging station operation with synthetic PBDR patterns of EV customers, and experimental results demonstrate that this framework can provide a more reliable prediction model for the ultimate optimization process, leading to more effective optimization solutions in terms of cost savings and charging station operation objectives with only a few training samples.

I. INTRODUCTION

Recent fast-growing penetration of electric vehicles (EVs) casts unforeseen challenges for charging station and power grid operations. Massive number of charging sessions inevitably result in overload due to the cluster effect on excessively concentrated EV charging stations and charging times [1], [2]. As both charging station and EV usually have limited capacities, excessive charging requests during peak times can cause considerable operational challenges. It is thus hard for charging stations to both satisfy consumers' demands and maximize charging profits.

Thus, it is desirable to alleviate the overload of the EV charging stations with appropriate charging control. And when optimizing the charging control policy, one non-negligible factor is the price-based demand response (PBDR) of EV chargers. Cost-sensitive EV customer can actively adjust their total amount of charging energy based on the assigned charging prices [3]. Therefore, in order to better operate simultaneous EV charging sessions, the price-based demand response (PBDR) strategy needs to be investigated through modeling the inherent relationship between energy

demands and price [4]–[8]. Knowing such PBDR patterns of EV customers, the charging station operator can precisely estimate customers' charging requirements and hence develop reasonable station-level charging policy to minimize operation cost while satisfying charging demands.

Previous research attempts utilized price-based demand response to encourage EV charging flexibility [9]. [5] turned to a rule-based approach to design the deadline-differentiated electric power service. Yet such rule-based methods heavily rely on a set of pre-defined heuristics or pre-assumed conditions. An on-off EV demand response strategy was leveraged by formulating a binary optimization problem [7], and finite-horizon Markov decision process was used to coordinate consumer owned energy storage [8]. These methods however simplified the demand response model such that *charging demands were assumed known exactly beforehand*. Moreover, both rule-based and optimization-based methods do not take into account the modeling procedure of price-based demand, and the charging demand is either estimated or assumed known exactly. An inaccurate forecast of price responsive charging load would further incur loss of robustness and reliability of station operation [10]. In this paper, we take a specific look into the relationship between charging price assigned to each EV owner and EV's charging demand, and investigate how to optimize charging strategy when such relationship is unrevealed to the charging station. We answer the following research question:

How do we optimize EV charging scheduling over unknown demands incurred by variable charging price?

We tackle both the demand forecasting and multi-step charging optimization problems in one framework. The proposed paradigm holds the promise of adaptively improving its performance over time when exposed to more data and feedback, making it well-suited for handling complex and dynamical problems. Standard pipelines involve individually fitting of a demand response function using the collected EV charging data, followed by the incorporation of predicted results into the ultimate optimization task [11], [12]. Such workflow neglects the fact that the locally accurate predictions may not benefit the underlying optimization problem's overall objectives [10], [13]. Moreover, the complex physical constraints make such predict-then-optimize framework produce infeasible solutions. [14] considered demand response model in an end-to-end fashion, while the application was limited to identifying a subset of demand response model parameters rather than solving the optimization problem.

Our formulation makes it possible to backpropagate from the ultimate charging task-based objective to the learning

¹ AI Thrust, Information Hub, The Hong Kong University of Science and Technology (Guangzhou). {cgu893, rliu519}@connect.hkust-gz.edu.cn, yizechen@ust.hk

² Division of Emerging Interdisciplinary Area under IPO, The Hong Kong University of Science and Technology. yuxin.pan@connect.ust.hk

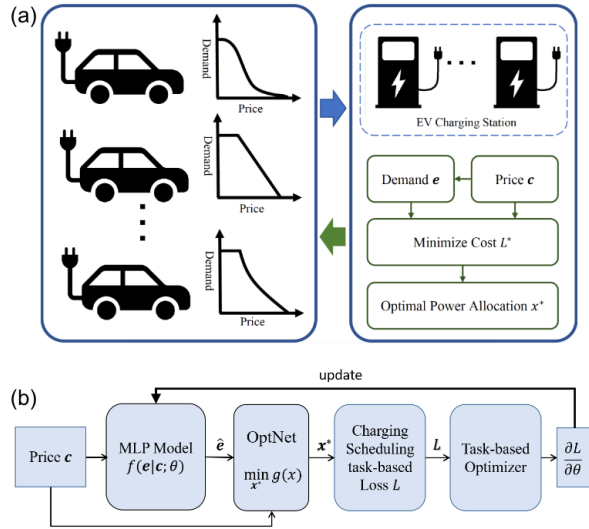


Fig. 1: (1a). Proposed EV demand estimation and charging station operation model. EV customers have various patterns of demand response to price. (1b). Proposed end-to-end charging demand forecasts and charging scheduling framework. With optimization task-based loss $\mathcal{L}$, our model propagates loss derivative with respect to neural network model's parameter θ to help train the prediction model $f(\cdot)$.

of individual PBDR model. Moreover, it elicits the control actions for EV charging. The schematic of proposed method is shown in Fig. 1. Specifically, to find the price-responsive EV charging demands, we learn the complex price-based demand response (PBDR) as a function of charging price unrevealed to the system operator. We combine prediction errors and downstream optimization cost errors in the model learning stage. Experimentally, our method is evaluated on noisy demand response samples based on varying charging prices. The results demonstrate that our framework can lead to a considerable decrease in the operation costs by more than 20% compared to the standard two-step (predict-then-optimize) method. This advantage is even more significant when the training dataset is small, meaning only limited customer responses are revealed to the charging stations.

II. PROBLEM FORMULATION

In this section, we describe the formulation of charging station operation problem, in which we further develop the relationship between the electricity selling price assigned to EV customers and their corresponding total electricity demand as a price-based demand response (PBDR) model.

A. Charging Station Operation Model

We first formulate the charging station operation model. We suppose the goal of one EV charging station is to find the optimal policy to minimize its operation cost, meanwhile satisfying its customers' charging demands as much as possible. Each EV comes to the charging station at various timesteps, and reveals their requested energy based on the charging price (costs in per kWh) along with the departure time. The charging station can control the charging rate of




Fig. 2: Examples of EV PBDR patterns. Individual EV's charging demand is a function of charging price.

each EV at each timestep. Denote N as the number of active EV customers at the charging station, $1 \leq t \leq T$ as the station's opening time horizon. Then the charging station scheduling problem is as follows:

$$\min_{\mathbf{x}} g(\mathbf{x}) = \sum_t p(t) \sum_i x_i(t) - \sum_i c_i \sum_t x_i(t) + \beta \left(\sum_i \left(\sum_t x_i(t) - e_i \right)^2 \right) + \alpha \sum_{i,t} x_i^2(t) \quad (1a)$$

$$\text{s.t.} \quad 0 \leq \sum_i x_i(t) \leq \bar{x}, \forall t \quad (1b)$$

$$0 \leq x_i(t) \leq \bar{y}_i, \forall i, t \quad (1c)$$

$$x_i(t) = 0, \text{ for } t \leq t_{a,i} \text{ and } t \geq t_{d,i} \quad (1d)$$

$$e_i = f_i(c_i); \quad (1e)$$

where $\mathbf{x} = [x_1(1), x_1(2), \dots, x_1(T), x_2(1), \dots, x_N(T)]^T$ is the control variable; $x_i(t)$ represents the charging power assigned to EV customer i at time t ; $p(t)$ is the electricity purchase price from the grid; c_i and e_i are the electricity selling price and total electricity demand for EV customer i . $\bar{x}$ is maximum charging power at time t for the station. $\bar{y}_i$ is the maximum charging power for EV customer i . With arrival time $t_{a,i}$ and departure time $t_{d,i}$, and β, α is the coefficient for regularizing demand charging completion and non-smooth charging power.

With known e_i , this EV scheduling problem is a constrained Quadratic Programming (QP) problem which strong duality holds, and it can be solved to find optimal charging policy. However, it can be challenging to solve in practice with unknown e_i . Since each EV owner has distinct charging behaviors, while the charging demand can be varying with respect to price. From the charging station side, it cannot solve (1a) without knowing such charging demand information. Customers adjust their electricity demands according to the selling price signal, and each customer may have its unique response to the price e_i . In order to model this price-demand map and use it to infer demand given price, we form the following demand price-response models. We note that in our setting, charging station operator informs customer c_i before charging starts, and the optimization of c_i along with charging scheduling is left for future work.

B. Price-based Demand Response Model

In this work we include two examples on the explicit modeling of price-based EV demand response which is not considered in previous research. We note that our solution

techniques described in Section III are not limited to the particular function mapping from charging price to charging demand, and we are able to solve the charging station optimization problem under various unknown price-responsive demand patterns.

Example 1: Utility-maximization EV demand

Once the EV owner i comes to the charging station and gets the price c_i revealed, the EV aims at minimizing the charging costs while maximizing the charging utility, given the prevailing electricity price. The utility function $f(e_i)$ is usually assumed to be continuous, nondecreasing, and concave [15]. In this paper, the user utility function is modeled as $f(e_i) = -(e_{level,i} - e_{trip,i})^2$, where $e_{level,i} = e_{init,i} + e_i \eta_i$. e_i , $e_{level,i}$, $e_{trip,i}$, and $e_{init,i}$ denote charging demand, final energy level, desired energy level after charging session, and initial energy level, respectively. Given that the users know their energy requirement for an upcoming trip, their utility function represents how well the charging need is met. Such utility maximization problem is formulated as

$$\min_{e_i} \quad \gamma_1 c_i e_i + \gamma_2 (e_{level,i} - e_{trip,i})^2 \quad (2a)$$

$$\text{s.t.} \quad e_{level,i} = e_{init,i} + e_i \eta_i \quad (2b)$$

$$SOC_{min,i} \leq e_{level,i} \leq SOC_{max,i} \quad (2c)$$

$$\frac{e_i}{(t_{d,i} - t_{a,i})} \leq p_{max,i} \quad (2d)$$

In the formulation above, Constraint (2b) states that the energy level $e_{level,i}$ of the i^{th} electric vehicle (EV) is equal to the sum of its initial energy level $e_{init,i}$ and the charging demand e_i ; (2c) requires that the energy level $e_{level,i}$ must fall within the battery state-of-charge capacity range; (2d) specifies that the average charging power rate must not exceed the maximum power rate. γ_1 and γ_2 denote the weighting factors of the charging cost and utility term; $\eta_i \in (0, 1)$ represents the charging efficiency. Generally, the charging demand e_i obtained by solving (2a) will have a piece-wise linear relationship with price signal c_i . Fig. (2)(a) illustrates one example of such pattern.

Example 2: Piecewise quadratic EV demand

Some EV owners are price-sensitive over certain price range, meaning that an increase in the charging cost results in a rapid decline in their demand for charging. To approximate the behaviour of this kind of customers, we adopt a general price-demand model and construct the piecewise function with quadratic components for their price-response model:

$$e_i = \begin{cases} e_{i,base}, & \text{if } c_i \leq c_{i,s} \\ \max(e_{i,base} - \mu_i (c_i - c_{i,s})^2, 0), & \text{if } c_i > c_{i,s} \end{cases} \quad (3)$$

where $e_{i,base}$ denotes the original charging demand of the i^{th} EV; μ_i represents the price sensitive parameter; when the current price c_i is greater than its sensitive price level $c_{i,s}$, the EV user decides to decrease charging demand. Fig. (2)(b) illustrates one example of demand price-response model based on the definition mentioned above.

III. END-TO-END LEARNING AND OPTIMIZATION

In practice, charging station is faced with unknown price-based demand response model for each customer. It is

challenging to design charging control under such situation. In this section, we illustrate how it is possible to include the training of price-demand forecasting module as a fully differentiable part in the EV charging optimization model.

A. Reformulation of Charging Station Model

To simplify the analysis, we first reformulate the optimization problem, and rewrite the charging station optimization problem (1a) as a standard Quadratic Programming (QP):

$$\begin{aligned} \min_{\mathbf{x}} g(\mathbf{x}) &= \beta (\mathbf{B}\mathbf{x} - \hat{\mathbf{e}})^T (\mathbf{B}\mathbf{x} - \hat{\mathbf{e}}) + (\mathbf{p}^T - \mathbf{c}^T \mathbf{B})\mathbf{x} + \alpha \mathbf{x}^T \mathbf{x} \\ &= \frac{1}{2} \mathbf{x}^T (2\beta \mathbf{B}^T \mathbf{B} + 2\alpha \mathbf{I}) \mathbf{x} + (\mathbf{p}^T - \mathbf{c}^T \mathbf{B} - 2\beta \hat{\mathbf{e}}^T \mathbf{B}) \mathbf{x} \\ &\quad + \beta \hat{\mathbf{e}}^T \hat{\mathbf{e}} \end{aligned} \quad (4a)$$

$$\text{s.t.} \quad \begin{bmatrix} -\mathbf{A} \\ \mathbf{A} \\ -\mathbf{I} \\ \mathbf{I} \end{bmatrix} \mathbf{x} \leq \begin{bmatrix} \mathbf{0} \\ \mathbf{r} \\ \mathbf{0} \\ \mathbf{y} \end{bmatrix} \quad (4b)$$

$$\mathbf{F}\mathbf{x} = \mathbf{0}. \quad (4c)$$

- $\mathbf{p} = [p(1), p(2), \dots, p(T), \dots, p(1), p(2), \dots, p(T)]^T \in R^{NT \times 1}$. lists the electricity price for charging $x_i(t)$.
- $\mathbf{c} = [c_1, c_2, \dots, c_N]^T$ lists the selling price assigned to each EV customer i .
- $\hat{\mathbf{e}} = [\hat{e}_1, \hat{e}_2, \dots, \hat{e}_N]^T$ lists the estimated demand of each EV customer i by a Multi-layer Perceptron (MLP) model $\hat{\mathbf{e}} = f(\mathbf{c}; \theta)$.
- $\mathbf{r} = \bar{x} \mathbf{1}_{(T \times 1)} \in R^{T \times 1}$ represents the charging capacity (i.e. station level maximum charging power).
- $\mathbf{y} = [\bar{y}_1 \mathbf{1}_{(T \times 1)}^T, \bar{y}_2 \mathbf{1}_{(T \times 1)}^T, \dots, \bar{y}_N \mathbf{1}_{(T \times 1)}^T]^T \in R^{NT \times 1}$ represents the customer level maximum charging power.
- $\mathbf{A} = [\mathbf{I}_{(T \times T)}, \dots, \mathbf{I}_{(T \times T)}] \in R^{T \times NT}$ helps transform term $\sum_i x_i(t)$ into matrix form. $\sum_i x_i(t) = \mathbf{A}\mathbf{x}$.
- $\mathbf{B} = \begin{bmatrix} \mathbf{1}_{T \times 1}^T & 0 & \dots & 0 \\ 0 & \mathbf{1}_{T \times 1}^T & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & \mathbf{1}_{T \times 1}^T \end{bmatrix} \in R^{N \times NT}$ helps transform term $\sum_i (\sum_t x_i(t) - \hat{e}_i)^2$ into matrix form. $\sum_i (\sum_t x_i(t) - \hat{e}_i)^2 = \mathbf{B}\mathbf{x} - \hat{\mathbf{e}}$.
- $\mathbf{F} \in R^{H \times NT}$ transforms the arrival and departure constraints (1d) into matrix form $\mathbf{F}\mathbf{x} = \mathbf{0}$. H is the total number of elements $x_i(t)$ which have constraint $x_i(t) = 0$. Each line F_h of $\mathbf{F}$ corresponds to one particular equality constraint $x_{i_h}(t_h) = 0$; in this line only element $F_{h, i_h \times T + t_h} = 1$.

Previous works usually suppose that if we have an accurate estimation $\hat{\mathbf{e}}$ of actual demand $\mathbf{e}$, the system cost (4a) will also achieve its optimal [16]. Hence, two-step approach separates prediction and optimization: first to train the MLP $f(\mathbf{c}; \theta)$ by minimizing root mean squared error (RMSE) between customer demand prediction and actual charging demand; then to use the trained MLP to predict customer demands $\hat{\mathbf{e}}$, and based on the predicted demands they solve problem (4) to obtain minimum operation cost.

Our goal is to learn the optimal parameter θ of an MLP prediction model $\hat{\mathbf{e}} = f(\mathbf{c}; \theta)$, which approximates

customer's demand response to price signal. Based on it we can minimize **actual** charging station operation cost $\mathcal{L}(\theta)$:

$$\min_{\theta} \mathcal{L}(\theta) = \beta(B\mathbf{x}^*(\theta) - \mathbf{e})^T(B\mathbf{x}^*(\theta) - \mathbf{e}) + (\mathbf{p}^T - \mathbf{c}^T B)\mathbf{x}^*(\theta) + \alpha\mathbf{x}^*(\theta)^T\mathbf{x}^*(\theta); \quad (5)$$

in which $\mathbf{x}^*(\theta)$ represents the optimal solution of the charging station operation optimization problem (4). $\mathbf{e}$ denotes the actual demand of each EV customer.

In contrast to the existing paradigm, our learning model is decision-focused, which means in the learning process we focus on the downstream optimization performance rather than solely on the prediction accuracy. The overall algorithm is an end-to-end framework to find the optimal decisions with price signals as input. Fig. (1) illustrates the architecture of our framework. More specifically, we directly use the objective function (5) of the charging station operation model as the training loss for the demand forecasting model MLP $f(\cdot)$, rather than the RMSE of demand prediction, and back propagate its gradient with respect to MLP model parameters $\frac{\partial \mathcal{L}}{\partial \theta}$, to adapt the model: $\frac{\partial \mathcal{L}}{\partial \theta} = \frac{\partial \mathcal{L}}{\partial \mathbf{x}^*} \frac{\partial \mathbf{x}^*}{\partial \theta}$. From (5) we can easily obtain $\frac{\partial \mathcal{L}}{\partial \mathbf{x}^*}$. Hence, we focus on computing the partial derivative $\frac{\partial \mathbf{x}^*}{\partial \theta}$ hereafter.

B. Differentiating the Optimal Solution

Now we analyze the optimal solution $\mathbf{x}^*$ of this standard form QP problem (4). Since strong duality holds, $\mathbf{x}^*$ should satisfy the Karush-Kuhn-Tucker (KKT) Conditions:

$$(2\beta B^T B + 2\alpha I)\mathbf{x}^* + (\mathbf{p}^T - \mathbf{c}^T B - 2\beta \hat{\mathbf{e}}^T B)^T + \begin{bmatrix} -A \\ A \\ -I \\ I \end{bmatrix}^T \boldsymbol{\lambda}^* + F^T \boldsymbol{\nu}^* = 0 \quad (6a)$$

$$F\mathbf{x}^* = 0 \quad (6b)$$

$$\text{diag}(\boldsymbol{\lambda}^*) \left(\begin{bmatrix} -A \\ A \\ -I \\ I \end{bmatrix} \mathbf{x}^* - \begin{bmatrix} \mathbf{0} \\ \mathbf{r} \\ \mathbf{0} \\ \mathbf{y} \end{bmatrix} \right) = 0 \quad (6c)$$

where $\boldsymbol{\lambda}^*$ and $\boldsymbol{\nu}^*$ are optimal dual variables. We compute the target partial derivative $\frac{\partial \mathbf{x}^*}{\partial \theta}$ by differentiating Equation (6). The matrix form of the KKT differentials is:

$$\underbrace{\begin{bmatrix} 2\beta B^T B + 2\alpha I & [-A^T A^T - I \quad I] & F^T \\ \text{diag}(\boldsymbol{\lambda}^*) \begin{bmatrix} -A \\ A \\ -I \\ I \end{bmatrix} & \text{diag}(\begin{bmatrix} -A \\ A \\ -I \\ I \end{bmatrix} \mathbf{x}^* - \begin{bmatrix} \mathbf{0} \\ \mathbf{r} \\ \mathbf{0} \\ \mathbf{y} \end{bmatrix}) & 0 \\ F & 0 & 0 \end{bmatrix}}_G \begin{bmatrix} \frac{d\mathbf{x}^*}{d\theta} \\ \frac{d\boldsymbol{\lambda}^*}{d\theta} \\ \frac{d\boldsymbol{\nu}^*}{d\theta} \end{bmatrix} = \begin{bmatrix} 2\beta B^T d\hat{\mathbf{e}} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \quad (7)$$

We divide both sides of Equation 7 with differential $d\theta$. And then we have a set of linear equations to solve for target derivative $\frac{\partial \mathbf{x}^*}{\partial \theta}$:

$$G \begin{bmatrix} \frac{d\mathbf{x}^*}{d\theta} \\ \frac{d\boldsymbol{\lambda}^*}{d\theta} \\ \frac{d\boldsymbol{\nu}^*}{d\theta} \end{bmatrix} = \begin{bmatrix} 2\beta B^T \frac{d\hat{\mathbf{e}}}{d\theta} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \quad (8)$$

Matrix G is obviously an upper triangular matrix here. Hence, it will always be invertible. Linear Equation 8 will achieve the unique solutions of gradients $\frac{d\mathbf{x}^*}{d\theta}$, $\frac{d\boldsymbol{\lambda}^*}{d\theta}$, $\frac{d\boldsymbol{\nu}^*}{d\theta}$. This gives the guarantee to make the gradients and parameter

updates to be well-defined. We refer to [17] for a more detailed discussion on the model differentiability.

C. Model Learning

In practice, we can solve QP problem (4) in batches efficiently in deep learning environment like Pytorch. We adapt the neural network architecture defined in OptNet [13]. Our approach is inspired by such end-to-end differentiable optimization framework [10], [13].

In our model, we output optimal solution $\mathbf{x}^*$ and compute the charging scheduling task-based loss gradient $\frac{\partial \mathcal{L}}{\partial \theta}$ by solving Equation (8). Then we backpropagate the gradient to update parameter θ in MLP model $f(\hat{\mathbf{e}}|\mathbf{c}; \theta)$ until the loss converges to an optimal value. Once the model is trained, we use it to generate estimated demand $\hat{\mathbf{e}}$. Then we forward $\hat{\mathbf{e}}$ to the differentiable optimization module, and it will solve (4) and output optimal solution $\mathbf{x}^*$. Given $\mathbf{x}^*$ we use (5) to compute the actual optimal operation cost $\mathcal{L}$. The whole end-to-end model learning process is illustrated in Algorithm. 1.

Algorithm 1 Charging Scheduling task based loss optimization

- 1: **Input:** price signal $\mathbf{c}$
 - 2: **Initialize** parameter θ for MLP model $f(\hat{\mathbf{e}}|\mathbf{c}; \theta)$.
 - 3: **for** Iteration j ($j = 1, 2, \dots, J_{max}$) **do**
 - 4: Estimate $\hat{\mathbf{e}} = f(\mathbf{c}|\theta)$.
 - 5: Solve Equation (4) for optimal solution $\mathbf{x}^*$ with estimated $\hat{\mathbf{e}}$.
 - 6: With $\mathbf{x}^*$ compute derivative $\frac{\partial \mathcal{L}(\mathbf{x}^*)}{\partial \theta}$ by solving Equation (8).
 - 7: Update $\theta_{j+1} \leftarrow \theta_j - \eta \frac{\partial \mathcal{L}}{\partial \mathbf{x}^*} \frac{\partial \mathbf{x}^*}{\partial \theta}$.
 - 8: **end for**
-

IV. NUMERICAL SIMULATIONS

In the simulation, we consider running a real-world charging station with electricity price data in Anaheim City, California [18], which has limited charging capacity and serves multiple EV customers for each day. We evaluate the performance with the overall actual optimal operation cost $\mathcal{L}(\theta)$, and 3 sub evaluation metrics: (1) *operation profit* $(\mathbf{p}^T - \mathbf{c}^T B)\mathbf{x}^*$, representing the profits earned by charging station. (2) *charging demand completion penalty* $\beta(B\mathbf{x}^* - \hat{\mathbf{e}})^T(B\mathbf{x}^* - \hat{\mathbf{e}})$, reflecting the demand satisfaction rate of EV customers. Lower penalty value represents higher demand satisfaction rate. (3) *smooth penalty* $\alpha(\mathbf{x}^*)^T\mathbf{x}^*$, reflecting the smoothness of charging curves under charging policy $\mathbf{x}^*$. We compare our framework with ground-truth results (*i.e.* solve (4) for optimal solution $\mathbf{x}^*$ under no demand estimation errors) and the two-step approach (as compared baseline). All models are built and trained with NVIDIA RTX 2080 Ti and Intel(R) Xeon(R) Silver 4210 CPU.

In the charging station model, we assume that the operational horizon of station is $T = 24$ hours since EV charging follows daily commute patterns, and the number of EV customers $N = 5$. Each EV customer can charge at a particular station at different starting times, so each EV




Fig. 3: (3a) shows the synthetic price-based demand response (PBDR) patterns used in the simulation; (3b) shows the real-world hourly commercial electricity purchase price [18] with colors indicating different price-time ranges. We set this peak-valley purchase price as p in our simulation.




Fig. 4: (a) operation costs during training, results averaged over 5 runs; (b) comparisons of sub-evaluation metrics.

customer has a unique charging session. We assign 3 Level-2 charging ports with 6.6 kW maximum charging power and 2 Level-2 charging ports with 3.6 kW maximum charging power for all 5 EV customers. The penalty coefficient for charging demand completion is $\beta = 5$, and the penalty coefficient for smooth is $\alpha = 0.001$. We construct 3 Piece-wise Linear and 2 Piece-wise Quadratic PBDR patterns as the synthetic PBDR patterns of 5 EV customers in the simulation based on Section II.B’s models (see Fig. (3a)). We set all charging demands of EVs varying in range [5, 20] (kW) based on general charging cases in Caltech ACN Dataset [19]. We also simulate customers with varying price sensitivity parameters and price sensitivity ranges. Based on the synthetic PBDR patterns, we generate a training dataset of 1,000 samples and a testing dataset of 100 samples. For the selling price c , we test two settings, where we assign random and different prices to 5 different EV customers [20]. Our method works well under both settings, and the price range [0.2, 0.6] (\$/kWh) is based on latest market price in U.S. 2023 [21]. Then we use synthetic PBDR patterns to compute corresponding actual demand response e . In this dataset we treat c as feature, e as label. We use real-world electricity purchase price shown in Fig. (3b) for the optimization.

In the numerical simulation, we examine how different methods minimize the actual operation cost $\mathcal{L}$ of the established charging station operation model with settings in Section IV.A. To benchmark our method, we develop both two-step approach and our decision-focused end-to-end framework (decision-focused E2E). Same as the network architecture applied in [10], we employ a commonly-used

2-hidden-layer network with a residual connection between inputs and outputs as the MLP model in both two-step approach and our decision-focused end-to-end framework. We respectively use 200, 400, 600, 800 samples in training dataset generated in Section IV.B to train the MLP model. For both methods we train the models for 250 iterations. We evaluate the performance of two methods with the average operation cost, profit, penalties of demand completion and smooth on the 100-sample test dataset in Section IV.B. All the results are benchmarked against ground-truth setting with known individual EV’s charging price and demands.

TABLE I: Evaluation on operation cost $\mathcal{L}$ (\$) with different number of training samples. Ground-Truth value is -11.69.

Samples	Decision-focused E2E	Two-step	Improvement(\$)
200	3.39 ± 0.45	7.01 ± 1.44	3.62
400	-4.90 ± 0.45	-1.91 ± 0.25	2.99
600	-7.26 ± 0.24	-4.45 ± 0.59	2.81
800	-8.07 ± 1.02	-5.56 ± 0.31	2.51

Fig. (4a) illustrates the convergence of station’s operation costs during 250 training iterations. The converged costs for both methods still have a gap compared to ground truth, as limited samples result in imperfect demand estimations. But as we expected, our decision-focused end-to-end framework outperforms the two-step approach, for it always converges to a better actual operation cost by explicitly training the forecasting MLP with the differentiable loss.

TABLE II: Evaluation on sub-metric *Operation Profit*, *Charging Demand Completion Penalty*, and *Smooth Penalty* with varying training samples. (5 runs average)

Samples	Operation Profits		Completion Penalty		Smooth Penalty	
	E2E	Two-step	E2E	Two-step	E2E	Two-step
200	15.41	15.33	18.62	22.15	0.1851	0.1861
400	15.42	15.36	10.33	13.26	0.1858	0.1871
600	15.39	15.34	7.95	10.70	0.1850	0.1865
800	15.45	15.40	7.19	9.65	0.1865	0.1870

Table I records the results of operation cost comparison under different numbers of training samples (All results are average of 5 runs). Our decision-focused end-to-end framework has more reliable performance on minimizing overall operation cost $\mathcal{L}$. The improvement upon two-step approach is much more apparent with small training dataset, indicating our framework performs better under the case of limited history. This demonstrates the value of incorporating the training of demand forecasting models into the overall optimization problem. In Table II, we show the corresponding results of 3 sub evaluation metrics: *operation profit*, *charging demand completion penalty*, *smooth penalty*. We can see that *charging demand completion penalty* contributes most to our framework’s leading over the two-step approach, while improvements on *operation profit* and *smooth penalty* are relatively mild, as illustrated in Fig. (4b). The major reason is that our framework optimizes the operation cost $\mathcal{L}$ by improving MLP estimations of charging demands $\hat{e}$, and *charging demand completion penalty* is the only sub-evaluation metric that contains a second-order term of




Fig. 5: Comparison of two-step approach and our decision-focused end-to-end framework on station-level charging power (one example from test dataset).

$\beta \hat{e}^T \hat{e}$; meanwhile in the simulation we assign a charging demand completion penalty coefficient $\beta = 5$. This makes *charging demand completion penalty* the most influenced metric by our proposed framework. In our formulation, we can also conveniently tune the penalty coefficients α, β in optimization objectives in Equation (4a) so that the demand prediction model can emphasize on particular task (e.g., maximizing operation profits, smoothing charging curves). Proposed method can strike a balance between shaping the charging curves and optimizing the charging needs.

Fig. 4a shows operation cost decreasing process of two-step approach and our decision-focused end-to-end framework; while Fig. 4b shows our decision-focused end-to-end framework outperforms two-step approach much on *charging demand completion penalty*. In Fig. 5, we further look into the overall charging power delivered by the station. Compared to two-step approach, we can validate that our framework achieves lower total charging power, which is closer to ground-truth value. The major improvements occur in the peak-price period when both profit and completion rate metrics are crucial, and decision-makers need to find a balanced solution away from the boundaries to minimize the total cost. This result justifies that by taking the decision objective into forecasting modules, our framework can achieve higher session completion rate while consuming less total power with imperfect demand estimation, especially for periods that require a delicate balanced solution.

V. DISCUSSION AND CONCLUSION

In this paper, we propose an end-to-end learning and optimization framework for price-responsive EV charging session scheduling. In particular, we model the EV charging demands as unknown values in response to varying pricing signals for EV users. Simulation results validate that our proposed framework demonstrates superior performance compared to the existing two-step approach (prediction-based learning + optimizing on standalone prediction). Our advantage is particularly pronounced in cases with small training datasets, indicating the effectiveness of our framework in situations with limited historical charging records. In the future work, we will explore more efficient computation for KKT differentiation to enable the framework to tackle realistic charging problems with huge EV customers, and consider real-time implementation with stochastic EV arrivals. We

will also try to enhance the generalization and robustness of EV PBDR modeling by training with more realistic, noisy datasets. Furthermore, we will explore the co-optimization of time-varying pricing assignments and charging schedules for individual customers. We are also interested in expanding the proposed end-to-end learning and optimization framework to encompass a broader range of demand response programs.

REFERENCES

- [1] S. W. Hadley and A. A. Tsvetkova, "Potential impacts of plug-in hybrid electric vehicles on regional power generation," *The Electricity Journal*, vol. 22, no. 10, pp. 56–68, 2009.
- [2] S. Li, H. Xiong, and Y. Chen, "Diffcharge: Generating ev charging scenarios via a denoising diffusion model," *IEEE Transactions on Smart Grid*, 2024.
- [3] M. H. Cedillo, H. Sun, J. Jiang, and Y. Cao, "Dynamic pricing and control for ev charging stations with solar generation," *Applied Energy*, vol. 326, p. 119920, 2022.
- [4] S. Shao, M. Pipattanasomporn, and S. Rahman, "Demand response as a load shaping tool in an intelligent grid with electric vehicles," *IEEE Transactions on Smart Grid*, vol. 2, no. 4, pp. 624–631, 2011.
- [5] E. Bitar and Y. Xu, "Deadline differentiated pricing of deferrable electric loads," *IEEE Transactions on Smart Grid*, vol. 8, no. 1, pp. 13–25, 2016.
- [6] A. J. Conejo, J. M. Morales, and L. Baringo, "Real-time demand response model," *IEEE Transactions on Smart Grid*, vol. 1, no. 3, pp. 236–242, 2010.
- [7] L. Yao, W. H. Lim, and T. S. Tsai, "A real-time charging scheme for demand response in electric vehicle parking station," *IEEE Transactions on Smart Grid*, vol. 8, no. 1, pp. 52–62, 2016.
- [8] Y. Xu and L. Tong, "On the operation and value of storage in consumer demand response," in *53rd IEEE Conference on Decision and Control*. IEEE, 2014, pp. 205–210.
- [9] C. Develder, N. Sadeghianpourhamami, M. Strobbe, and N. Refa, "Quantifying flexibility in ev charging as dr potential: Analysis of two real-world data sets," in *2016 IEEE International Conference on Smart Grid Communications*. IEEE, 2016, pp. 600–605.
- [10] P. Donti, B. Amos, and J. Z. Kolter, "Task-based end-to-end model learning in stochastic optimization," *Advances in neural information processing systems*, vol. 30, 2017.
- [11] R. T. Rockafellar and R. J.-B. Wets, "Scenarios and policy aggregation in optimization under uncertainty," *Mathematics of operations research*, vol. 16, no. 1, pp. 119–147, 1991.
- [12] I. Antonopoulos, V. Robu, B. Couraud, and D. Flynn, "Data-driven modelling of energy demand response behaviour based on a large-scale residential trial," *Energy and AI*, vol. 4, p. 100071, 2021.
- [13] B. Amos and J. Z. Kolter, "OptNet: Differentiable optimization as a layer in neural networks," in *Proceedings of the 34th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 70. PMLR, 2017, pp. 136–145.
- [14] Y. Bian, N. Zheng, Y. Zheng, B. Xu, and Y. Shi, "Demand response model identification and behavior forecast with optnet: A gradient-based approach," in *Proceedings of the Thirteenth ACM International Conference on Future Energy Systems*, 2022, pp. 418–429.
- [15] N. Li, L. Gan, L. Chen, and S. H. Low, "An optimization-based demand response in radial distribution networks," in *2012 IEEE Globecom Workshops*. IEEE, 2012, pp. 1474–1479.
- [16] M. H. Amini, A. Kargarian, and O. Karabasoglu, "Arima-based decoupled time series forecasting of electric vehicle charging demand for stochastic power system operation," *Electric Power Systems Research*, vol. 140, pp. 378–390, 2016.
- [17] S. Barratt, "On the differentiability of the solution to convex optimization problems," *arXiv preprint arXiv:1804.05098*, 2018.
- [18] "Commercial rates. city of anaheim." <https://www.anaheim.net/6440/Commercial-Rates>, accessed: 2023-09-11.
- [19] Z. Lee, T. Li, and S. H. Low, "Acn-data: Analysis and applications of an open ev charging dataset," in *Proceedings of the Tenth International Conference on Future Energy Systems (e-Energy '19)*, June 2019.
- [20] "The shocking truth: How the price of a charging station varies based on the type of vehicle being charged." <https://energy5.com/>, accessed: 2023-09-27.
- [21] "How much does it cost to charge an electric car?" <https://www.enelxway.com/us/en/resources/blog>, accessed: 2023-09-11.