Applied Energy 402 (2026) 127032

Contents lists available at ScienceDirect

Applied Energy

journal homepage: www.elsevier.com/locate/apenergy

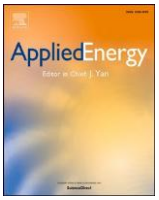

A novel differentiable predictive control (DPC) approach for safe and optimal EV charging and discharging scheduling

Yuewei Li [a], Bing Dong [a,*], Xuezheng Wang [a], Yueming Qiu [b]

[a] Department of Mechanical & Aerospace Engineering, Syracuse University, 263 Link Hall, Syracuse, NY 13244, United States
[b] School of Public Policy, University of Maryland at College Park, College Park, MD 20742, United States

HIGHLIGHTS

- Proposed a novel DPC approach for safety assurance in hard-constrained optimization problems.
- Introduced a projection layer derived from KKT conditions to enforce equality constraints within neural network-based control.
- Developed a novel barrier function method to effectively handle inequality constraints.
- Applied the framework to a multi-objective EV charging and discharging optimization problem, demonstrating practical effectiveness.
- Achieved significantly reduced constraint violations and improved objective performance compared to standard DPC approach.

ARTICLE INFO

Keywords:
Differentiable predictive control (DPC)
Deep neural networks
DPC formulation

ABSTRACT

Differentiable Predictive Control (DPC) has emerged as a data-driven alternative to Model Predictive Control (MPC), attracting growing interest across various research domains. Unlike traditional MPC, which relies on physics-based models, DPC utilizes black-box models—typically neural networks—to learn control policies. The primary advantage of DPC is its computational efficiency. However, a key challenge lies in ensuring constraint satisfaction, as neural network-based policies often struggle with maintaining feasibility. In conventional DPC frameworks for building-related studies, both equality and inequality constraints are enforced through penalty methods. However, tuning penalty parameters is complex, and constraint violations often persist. To address this issue, we propose a novel DPC approach that strictly enforces constraints. Specifically, we incorporate a projection layer based on Karush-Kuhn-Tucker (KKT) conditions to satisfy equality constraints. For inequality constraints, we replace the traditional penalty method with barrier functions, which provide steeper gradient as the solution approaches constraint boundaries, enforcing stricter constraint compliance. To evaluate the effectiveness of our approach, we apply it to an electric vehicle (EV) charging and discharging optimization problem, considering three objectives: minimizing electricity costs, reducing CO_2 emissions, and load flattening. We compare our method against conventional penalty-based DPC and the augmented Lagrangian method. The results indicated that the proposed approach reduced constraint violations by over 90 % in occurrence, compared to the conventional penalty method. Additionally, objective performance improved, achieving a 3 % greater bill reduction, 12 % more CO_2 reduction, and 29 % greater peak load reduction.

1. Introduction

1.1. What is DPC?

Model Predictive Control (MPC) has seen rapid development across various control applications. Despite advancements in intelligent control strategies, most buildings still rely on rule-based control (RBC) systems, which offer only limited energy savings [1–3]. These conventional approaches typically implement predefined operational sequences without accounting for the dynamic nature of building performance, weather conditions, or occupancy patterns. However, the decreasing cost of computational power and sensing technologies is making more sophisticated control approaches, such as MPC, increasingly viable. Over the past decade, MPC has emerged as a leading method in intelligent building operations due to its ability to

* Corresponding author.
E-mail address: bidong@syr.edu (B. Dong).

https://doi.org/10.1016/j.apenergy.2025.127032
Received 24 March 2025; Received in revised form 6 August 2025; Accepted 3 November 2025

Nomenclature

Abbreviations

MPC	model predictive control
RBC	rule-based control
LBMPC	learning-based MPC
eMPC	explicit MPC
aMPC	approximate MPC
sMPC	stochastic MPC
DPC	differentiable predictive control
AD	automatic differentiation
SGD	stochastic gradient descent
PINN	physics-informed neural network
KKT	Karush-Kuhn-Tucker
SOC	state of charge
NIOM	non-intrusive occupancy monitoring
LSTM	long short-term memory

Symbols

x	model input
y	model output
W	weight matrix
b	bias matrix
ω	coefficient
Var	variance
B	barrier function
μ	barrier parameter
ℓ	loss function
υ	augmented Lagrangian multiplier for 2nd order term
φ	augmented Lagrangian multiplier for 1st order term
η	battery charging and discharging efficiency
M, N	state space model parameter matrix
L	load
P	charging/discharging power
λ	electricity price
E	CO_2 emission marginal factor
O	occupancy presence
S	battery state of charge
η	battery charging/discharging efficiency

Subscripts

$\cdot_t$	time step
$\cdot_b$	electricity bill
$\cdot_c$	CO_2 emission
$\cdot_v$	load flatten
$\cdot_l$	constraint for load
$\cdot_p$	constraint for charging/discharging power
$\cdot_s$	constraint for battery state of charge
$\cdot_{s_f}$	constraint for battery state of charge when departure

Superscripts

$\cdot^i$	household index
$\cdot^*$	correction

Accents

$\hat{\cdot}$	projection

systematically optimize thermal comfort while achieving substantial energy savings—ranging from 15 % to 50 %, as demonstrated in various simulations and pilot studies [4–7]. Furthermore, MPC plays a key role in demand-side management by enabling price-responsive control and active participation in demand response programs [8–11]. By leveraging mathematical models to predict future building behavior, MPC determines optimal control actions that balance energy efficiency, occupant comfort, operational constraints, and external factors such as weather forecasts in a structured and adaptive manner.

Several variations of MPC have been developed to address different challenges and improve computational efficiency. Learning-Based MPC (LBMPC) integrates system identification techniques to model system dynamics from data, extending traditional adaptive MPC frameworks [12,13]. Explicit MPC (eMPC) precomputes optimal control laws offline, significantly reducing real-time computational requirements. However, eMPC suffers from the curse of dimensionality, where the complexity of precomputed solutions grows exponentially with the number of constraints, limiting its applicability to small-scale systems with short prediction horizons [14]. Other approaches, such as Approximate MPC (aMPC) and Stochastic MPC (sMPC), address uncertainties but require solving complex optimization problems at each time step, making real-time implementation challenging in large-scale systems with stringent computational constraints. As system complexity increases, MPC's computational burden and scalability issues become more pronounced, restricting its application in real-time, resource-constrained environments. Neural network (NN-) based MPC was then proposed. While both DPC and NN-based MPC leverage machine learning networks within control optimization, their roles and learning mechanisms are fundamentally different. In NN-based MPC, the network is used to approximate the system dynamics. However, the control action is still determined through an online optimization process at each time step. In contrast, DPC treats the entire control formulation as a differentiable computational process, allowing the control policy itself to be trained offline using gradient-based optimization [15].

To address these challenges, researchers have explored differentiable approaches to solving MPC problems, leading to the development of Differentiable Predictive Control (DPC). DPC provides a framework for learning explicit predictive control policies offline [15–19]. It leverages differentiable programming [20], where the mathematical operations defining the MPC problem—objective functions and constraints—are structured as a directed acyclic computational graph within an automatic differentiation (AD) framework. This enables direct computation of policy gradients, facilitating constrained policy learning through stochastic gradient descent (SGD). Unlike conventional MPC, which requires solving an optimization problem at every control step, DPC learns stable neural control policies in an end-to-end manner, making it a scalable alternative for constrained optimal control problems. DPC combines the predictive capabilities of MPC with the computational advantages of neural networks, creating a hybrid approach that is both effective and efficient. By formulating the MPC problem within a differentiable framework, DPC enables the direct training of neural control policies that can approximate the behavior of MPC controllers without the need for online optimization. Additionally, DPC eliminates the need for an explicitly calibrated physics-based or white-box model, further reducing computational complexity and simplifying model development. Due to these advantages, DPC has been widely applied in various fields, including building energy management [21], traffic network control [22], robust control in robotics [23], economic dispatch in power systems [24], and end-to-end planning and control [25].

1.2. Challenges for DPC

Despite its potential, DPC faces significant challenges when applied to complex optimization problems with multiple competing objectives and operational constraints. One major limitation is model tuning—balancing multiple objectives and constraints in the loss function requires careful weight adjustments to avoid suboptimal solutions, whereas the DPC performance is sensitive to cost function design and

weighting [19,26–28]. Moreover, DPC does not inherently guarantee strict constraint enforcement. Unlike traditional MPC, which explicitly satisfies constraints through optimization solvers, DPC reformulates constraints within the loss function, typically using penalty-based methods [15,17,19]. However, this approach only ensures soft constraint satisfaction, meaning large penalty coefficients can introduce numerical instability, while small coefficients may lead to constraint violations. For instance, in the open-source DPC implementation project [17] constraints are integrated into the loss function via penalty methods using mean absolute error (MAE) as the metric. While this formulation provides a differentiable training objective, it lacks a rigorous mechanism for strict constraint enforcement, which can be problematic for safety-critical applications where constraint violations could lead to system failure or safety hazards. For example, the case study involved in this study is for optimizing EV charging and discharging schedule. Despite the growing research on EV charging optimization, several key challenges remain unresolved. First, many existing approaches, including those based on MPC and rule-based heuristics [29,30], struggle to rigorously enforce hard operational constraints, violations of which can pose safety and grid reliability risks. Second, real-world EV scheduling problems are inherently multi-objective, involving trade-offs between generation cost and emissions [31]. Balancing these objectives dynamically across many users introduces complexity that traditional penalty-based methods handle difficultly. Third, user behaviors such as arrival time, departure SOC, and occupancy vary substantially, requiring control strategies that generalize across scenarios [32]. Moreover, many advanced control frameworks are computationally burdensome, limiting their feasibility for fast and large-scale deployment. These gaps motivate the development of a scalable and constraint-compliant control framework capable of learning feasible and efficient charging policies under realistic EV operational conditions. Regarding the constraints satisfaction, the limits for EV charging and discharging power are one of the hard inequality constraints that would cause safe-critical problems. The violations on EV charging and discharging power will also affect the overall integrated load profiles. Without self-generation, the energy consumption is critical to be non-negative, which is also a safety-critical problem. Only with all hard constraints complied, can the results and model performance be evaluated fairly.

To mitigate these issues, researchers have explored alternative constraint-handling techniques. In image segmentation applications, for example, the augmented Lagrangian approach has been used to improve equality constraint enforcement [33,34]. This method introduces Lagrange multipliers that adaptively adjust the penalty weights during training, leading to better constraint satisfaction without requiring excessive penalty coefficients. More broadly, the field of Physics-Informed Neural Networks (PINNs) has introduced novel methods for incorporating physical constraints into machine learning models. One such approach, the KKT-hPINN framework, which refers to a novel PINN that rigorously guarantees hard linear equality constraints, uses a projected layer derived from Karush-Kuhn-Tucker (KKT) conditions [35]. This projection ensures that the model's output remains within the feasible region without requiring equality constraints to be embedded in the loss function. For inequality constraints, researchers have proposed using logarithmic barrier functions as an alternative to traditional penalty methods [36,37]. These functions increase gradient magnitudes near constraint boundaries, guiding neural control policies to remain within feasible operating conditions more effectively.

1.3. Objectives

In this study, we propose a novel DPC framework that integrates the KKT-hPINN approach for equality constraint enforcement and a barrier function method for handling inequality constraints. By embedding these techniques into the DPC structure, we enhance constraint satisfaction while reducing the reliance on extensive model tuning. To demonstrate the effectiveness of our approach, we apply it to a community-level EV charging and discharging optimization problem. This domain was chosen due to the complexity and difficulties of precisely modeling EV behaviors, which poses a strong challenge for utilizing traditional physics-based solutions. Moreover, due to its combination of multi-objective requirements, hard operational constraints, and mixed constraint types, which poses a strong challenge for traditional DPC frameworks. Unlike synthetic benchmarks, EV scheduling inherently involves critical constraints such as battery SOC limits, charging power bounds, and departure readiness requirements, where these violations can compromise user experience and grid stability. Additionally, it demands simultaneous optimization of users' electricity bills, CO_2 emissions, and load flattening, making it a compelling test case for validating the proposed constraint-handling mechanisms. Through this case study, we illustrate the benefits of our enhanced DPC framework in improving constraint enforcement and optimization performance. The selection of EV charging optimization as our application domain is also motivated by the increasing integration of EVs into the power grid, which presents new challenges in balancing supply and demand. As electric vehicle adoption continues to grow, effective charging strategies become crucial for mitigating peak demand issues and improving grid stability. Uncoordinated EV charging can exacerbate peak demand periods, potentially leading to grid instability and increased infrastructure costs. Comparing to traditional rule-based charging strategies, MPC-based methods provide a more systematic approach to optimizing charging schedules by considering future electricity prices, grid conditions, and user requirements [38–40]. By leveraging DPC, we enable scalable and adaptive EV charging coordination, ensuring both user convenience and energy system efficiency while respecting the operational constraints of both the vehicles and the power grid.

To validate our approach, we compare the performance of four DPC formulations in model training: (a) the conventional penalty method, (b) the augmented Lagrangian (AL) method for equality constraints, (c) the KKT-based projection layer for equality constraints, and (d) our proposed novel method, which integrates the KKT-based projection with a barrier function for enforcing both equality and inequality constraints. Through a comprehensive simulation-based analysis, we demonstrate the effectiveness of our approach in enhancing constraint satisfaction while improving optimization performance across all objectives.

In this paper, we expand the capabilities of DPC with following contributions:

1) We propose a hybrid DPC framework that integrates a KKT-based projection layer for hard equality constraint enforcement and a modified barrier function for inequality constraints, improving constraint satisfaction without relying on extensive penalty tuning.
2) We demonstrate the effectiveness of the proposed framework on a multi-objective EV charging and discharging scheduling problem using real-world smart meter data, showing improved performance in constraint compliance and optimization performances compared to existing DPC formulations.
3) We show that the proposed structure simplifies training by reducing sensitivity to hyperparameter tuning and improving convergence, supporting its scalability to larger scale of study.

The remainder of this paper is structured as follows: Section 2 introduces the methodology for the improved DPC framework, including the neural network architecture and the proposed constraint handling techniques; Section 3 presents the case study and simulation setup, including the dataset description, data preprocessing, problem formulation, and DPC implementation details; Section 4 presents a comprehensive analysis of the results, comparing the performance of different DPC formulations across various metrics; Section 5 discusses the limitations and future work; and Section 6 concludes the study with key findings.

2. Methodology

This section introduces the methodology of the proposed DPC framework, including the structure and the approaches for handling equality and inequality constraints. Fig. 1 depicts those approaches in general. The proposed barrier function is a refined approach for managing inequality constraints within the loss function. Additionally, a projection layer, derived from KKT conditions, is integrated into the DPC framework to ensure compliance with equality constraints, as adopted from [35].

2.1. DPC structure

A fully connected neural network serves as the foundation of the DPC framework. The architecture follows a feed-forward structure, where the forward computation is expressed as:

$$z_{(l)} = \sigma\left(W_{(l-1)} z_{(l-1)} + b_{(l-1)}\right) \tag{1}$$

In this framework, $z_{(l)}$ represents the output of layer l. In that context, $z_{(0)}$ denotes the model input, while $z_{(L)}$ corresponds to the output of the final layer L. At each layer l, the associated outputs $z_{(l)}$ are computed based on the activations from the previous layer $z_{(l-1)}$, which are transformed through the weight matrix $W_{(l-1)}$, bias term $b_{(l-1)}$, and an activation function σ.

In DPC, the neural network operates within an unsupervised learning framework, meaning there is no explicit target data to compare with the final output of the network. Instead, the DPC network's goal is to generate optimized control sequences through the iterative training process. This process is driven by a loss function, which is carefully designed to incorporate both the objectives and constraints specific to the problem at hand. The network is trained to minimize this loss function, which balances the various competing factors involved in the control task. The training process for the DPC network can be formulated as follows:

$$\theta^* = argmin\, L(\theta) \tag{2}$$

where θ^* represents the optimized control actions, and $L(\theta)$ represents the loss function.

2.2. Safety assurance for hard equality constraint

Building on the work from [35], we introduced an orthogonal projection layer based on KKT conditions to map the final output to its closest feasible region. Fig. 2 illustrates the overall neural network structure incorporated with the projection layers. Specifically, two fully connected layers are integrated into the neural network as a correction mechanism, ensuring that the predictions are adjusted to the nearest feasible point. These correction layers are non-trainable, with fixed weights and biases determined from the KKT conditions, that are determined by the equations as following.

In the case of a hard linear equality constraint as $Ax + By = b$, where x represents the input and y is the original output variable, the nearest feasible solution $\tilde{y}$ can be expressed as follows:

$$\tilde{y} = A^* x + B^* y + b^* \tag{3}$$

Where A^*, B^* and b^* are the correcting matrixes, and can be expressed as:

$$A^* = -B^T\left(BB^T\right)^{-1}A \tag{3}$$

$$B^* = \mathbf{I} - B^T\left(BB^T\right)^{-1}B \tag{4}$$

$$b^* = B^T\left(BB^T\right)^{-1}b \tag{5}$$

The two non-trainable layers are added separately to guide the neural network outputs toward y. One layer is a linear transformer that processes the input x using fixed weights A^* and a fixed bias term b^*, while the other layer transforms the original output y using the fixed weights B^*. The projected output $\tilde{y}$ is then obtained by adding the outputs from these two distinct layers together. For a detailed theorem and proofs, we refer the readers to [35].

2.3. Safety assurance for hard inequality constraints

Unlike the safety assurance for equality constraints, which is achieved through the enhanced network structure, inequality constraints are ensured through an innovative model training process. As mentioned earlier, the control law in DPC is learned by updating the model parameters during training, based on a loss function derived from the control objectives and constraints. In this context, we employ a modified barrier function approach to replace the traditional linear penalty method for handling inequality constraints. This approach guarantees the strict satisfaction of these hard constraints, ensuring that the model respects the inequality conditions throughout the training process.

The traditional barrier function method transforms a constrained optimization problem into an unconstrained one by incorporating the inequality constraints into the objective function. This is achieved by introducing a penalty term that replaces the inequality constraints. Specifically, the penalty term is designed to penalize any violation of the constraints stricter, effectively guiding the optimization process to avoid infeasible solutions. The traditional logarithmic barrier functions are explained in **Eqs. (6)–(8)**. Consider the following constrained optimization problem:

$$\min_x f(x) \tag{6}$$

$$s.t. g(x) \le 0 \tag{7}$$

Using log barrier function, the previous constrained optimization problem can be reformulated as:

$$\min_x f(x) - \mu \log(-g(x)) \tag{8}$$

where μ is the barrier parameter, which is a positive scalar that controls the influence of barrier term. The reformulated unconstrained optimization problem is equivalent to the constrained problem as **Eq. (6) and**

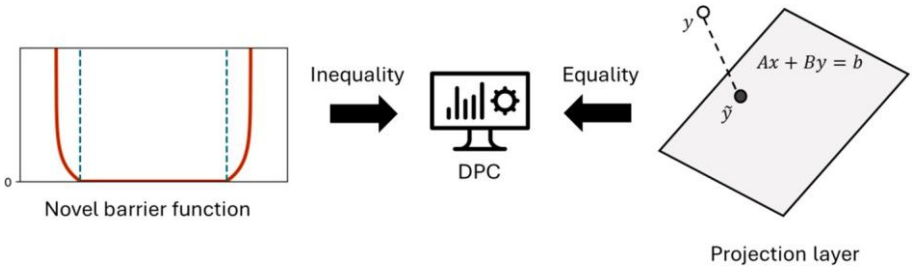

Fig. 1. Novel measures for handling inequality and equality constraints in the proposed DPC approach.

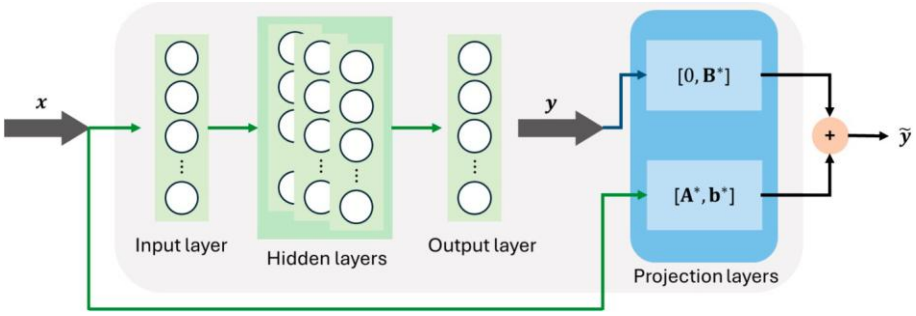

Fig. 2. Neural network structure (green blocks are trainable layers; blue blocks are the two non-trainable projection layers). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

(7). The log function penalizes any violation of the constraints, and as μ increases, the solution approaches feasibility with respect to the inequality constraints while ensuring that the inequality constraints are satisfied as the optimization progresses. The reformulated equation will exhibit low values within the feasible region, while the values increase toward infinity as the solution approaches the boundaries of the feasible region. As a result, during the training process, the optimization will push the solution toward the interior of the feasible region, where the penalty values are minimized.

However, the traditional barrier function is not suitable for our case study, where optimal solutions are not lying at the interior of the feasible region. In our case, one of the objectives is to minimize electricity bills, which could require EVs to discharge as much as possible during peak times. One of the inequality constraints in our study is the EV battery's charging and discharging power limitations. Because we want the battery to actively charge or discharge rather than remain idle at a neutral state (zero), the traditional barrier function is unsuitable as it would drive the optimized solution toward the midpoint of the feasible region.

To address this, we developed a novel version of the barrier function. In this approach, the loss values remain zero as long as the solution stays within the boundaries of the feasible region. Once the output exceeds the boundaries, the penalty increases sharply toward infinity. The proposed barrier functions are depicted in **Eqs.** (9)–(12). This modified barrier function ensures that the optimization process pushes the solution toward the desired extremes of the feasible region, as required by the problem. Consider an optimization problem formulated as follows:

$$\min_x f(x) \tag{9}$$

$$s.t. a \le g(x) \le b, a, b \in \forall \tag{10}$$

Using the proposed barrier functions, the problem can be reformulated as:

$$\min_x f(x) + \boldsymbol{B}(x) \tag{11}$$

$$\boldsymbol{B}(x) = \begin{cases} -\mu \log(1 - (g(x) - b)), g(x) > b \\ 0, a \le g(x) \le b \\ -\mu \log(1 - (a - g(x))), g(x) < a \end{cases} \tag{12}$$

Where $\boldsymbol{B}(x)$ represents the barrier function term for the inequality constraints. Fig. 3 illustrates behavior of the barrier terms for the EV power inequality constraints, with different barrier parameters μ selected. The modified barrier function ensures that the outputs within the boundaries of the feasible region receive the same penalty. Nonetheless, if the output goes beyond the boundaries, the steep curve results in a significant increase in the penalty, which drives the solution back toward the feasible region, ensuring that the inequality constraints are satisfied. It can be observed that as μ increases, the curve becomes steeper, meaning the penalty's influence on the network outputs grows stronger, providing a more robust safety assurance. This modified barrier function approach was applied to all inequality constraints in the model.

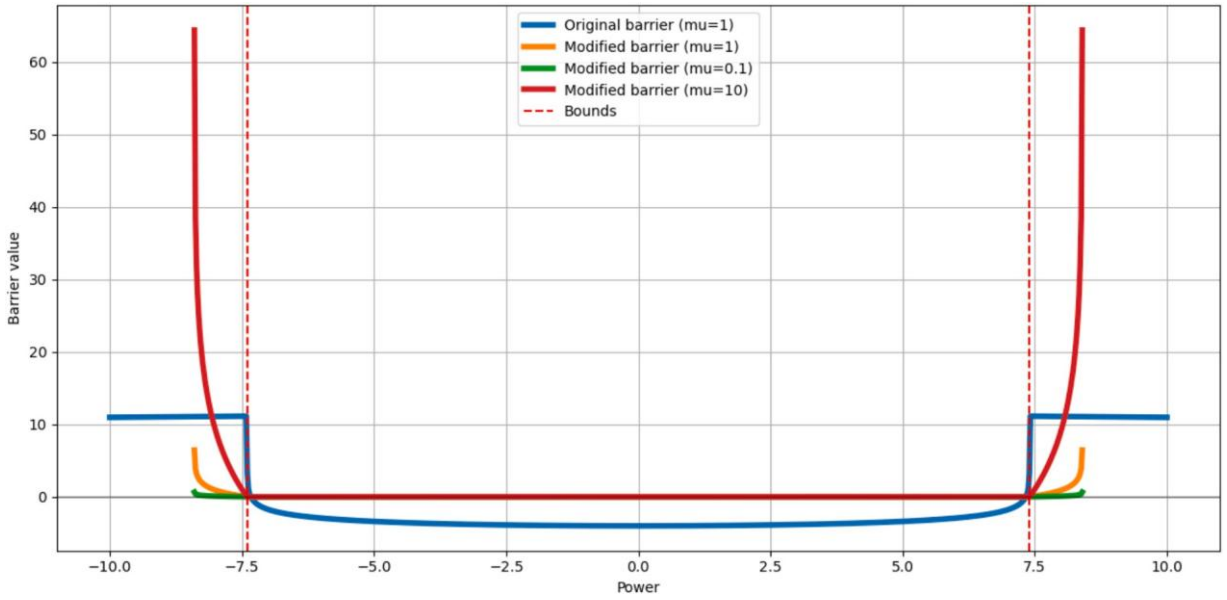

Fig. 3. Modified barrier function behaviors with different barrier parameters.

3. Case study

This section introduces the case study for optimizing EV charging and discharging schedules. Fig. 4 depicts the technical roadmap of the case study. Description of the dataset, data preprocessing, problem formulation, and different DPC models are explained in detail as follows.

3.1. Dataset description

The smart meter dataset is provided by a major utility company serving the metropolitan area of a large city in the southwestern United States. The dataset includes hourly smart meter data for individual households from 2013 to 2019. For this study, we specifically selected 84 users who enrolled in rate plan 29, a time-of-use (TOU) plan designed for EV users. This plan includes on-peak, off-peak, and super off-peak pricing, with the details presented in Fig. 5.

We focused on households that exclusively adopted EVs, excluding those with additional distributed energy resources (DERs). For these EV-only customers, we also had access to another dataset containing their EV adoption dates. This information was used to separate the raw smart meter data into two parts: before and after EV adoption. By comparing these two segments, the charging power for each household can be empirically determined.

This study focused solely on weekday scenarios, as customers are more likely to commute between home and workplace during these days. Instead of using a calendar day, a daily scenario was defined from 12 PM to 12 PM the following day (with EV charging typically occurring at night, before occupants leave for work). This time window alignment provides a more realistic representation of how EVs are typically used in residential settings, where vehicles are often away during working hours and available for charging overnight. The dataset consisted of 22 weekdays from August 2018, with the first 21 days used for training and the final day reserved for testing. It is important to note that the testing data includes 2016 hourly data points (84 households $\times$ 24 h), resulting in a total control sequence size of 2016.

3.2. Data preprocessing

The data preprocessing process involves two major steps. The first step is occupancy detection, which helps to better understand the EV users' occupancy schedules. However, the smart meter data only contained hourly smart meter readings without detailed miscellaneous loads, creating a challenge for directly correlating occupancy presence. To address this, we adopted a non-intrusive occupancy monitoring (NIOM) approach, which uses three different metrics to infer occupancy through a threshold-based statistical method [41]. These metrics include average power, standard deviation, and power range, as occupied homes typically exhibit higher energy usage and greater variability in power consumption compared to unoccupied ones. Moreover, occupancy is associated with a broader range of power consumption. The NIOM algorithm ran dynamically for each daily scenario and was exclusively used for detecting daytime occupancy (7 AM to 11 PM). For nighttime occupancy (12 AM to 6 AM), we inferred occupancy from the preceding evening, assuming that normal nighttime loads remain unaffected by occupancy presence. This served as a rough estimate of energy usage during unoccupied times, with thresholds also derived from this assumption. This approach worked well for typical workday scenarios, given the high likelihood that individuals stay at home during nighttime hours unless they are out of town or working night shifts. A key benefit of this method is that it does not require ground truth home occupancy data for model training, thereby reducing the need for data collection and addressing potential privacy concerns.

The second step involves estimating the baseline building load and the arrival State of Charge (SOC), both of which are crucial input features for the DPC network. Since the raw dataset only contained smart meter readings, we combined the baseline building load and original EV charging power by developing a Long Short-Term Memory (LSTM) model to separate EV charging loads from the building's energy usage. We generated labeled data by adding random EV charging profiles to pre-adoption smart meter data for model training and verification. After training, the model was applied to post-adoption data to decompose the charging loads. This model also estimated the total daily charging load, which allowed us to backtrack and calculate the arrival SOC, under the assumption that the EV leaves home at the maximum SOC as defined.

3.3. Problem formulation

For this case study, we formulated a multi-objective optimization problem aimed at optimizing the scheduling of EV charging and discharging behaviors. The objectives are to minimize the users' overall electricity bills, reduce CO_2 emissions, and flatten the aggregated load profiles of the total I households as much as possible. The objective

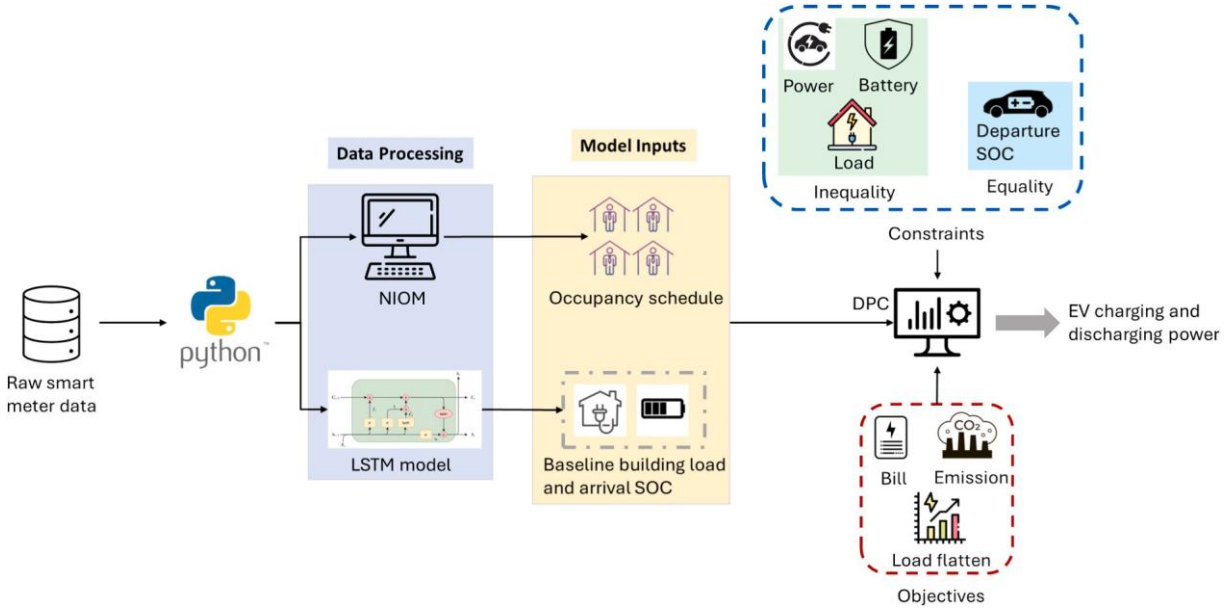

Fig. 4. Technical roadmap for optimal EV charging/discharging scheduling.

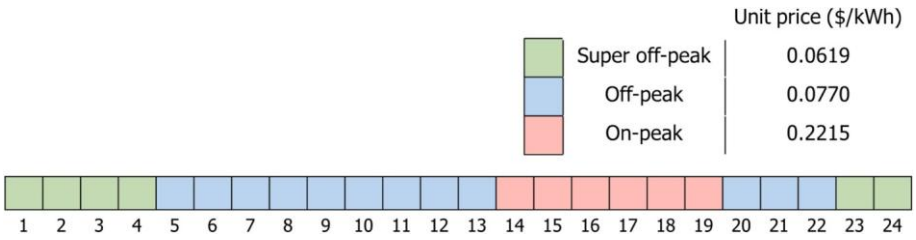

Fig. 5. Time-of-use (TOU) price plan.

function can be expressed as:

$$\min_{P_t^i} \sum_I \sum_T \omega_i \, L_t^i\left(P_t^i\right) \lambda_t^i + \omega_c \left[E_t \left(L_t^i\left(P_t^i\right) - L_{t,m}^i \right) \right] + \omega_p \, Var\left(L_t^i\right) \tag{13}$$

where P_t^i is the EV charging and discharging power for household i at time t, which is the control variable. The positive values of P_t^i represent charging power and negative values represent discharging power. $L_t^i\left(P_t^i\right)$ represents the load variable of household i at time t, which is a function of the optimal variable. λ_t^i is the electricity price for household i at time t, which is adopted from the rate plan that the user is enrolled in. E_t is the marginal CO_2 emission factor at time t, which are summarized in Table 1. $L_{t,m}^i$ is the measured smart meter reading. $L_{t,\beta}^i$ represents the 24-h load pattern for household i. The first term of **Eq. (13)** represents the utility bill, calculated as the product of the electricity price and energy usage. The second term accounts for the reduction of CO_2 emissions, which is quantified using the marginal emission factors and the load differences. The final term represents the variance of the aggregated load profiles, which is minimized to flatten the load pattern. ω_b, ω_c and ω_p are the weights arranged for each of the objectives, respectively.

The constraints of the optimization problem are as follows:

$$L_t^i\left(P_t^i\right) = L_{t,\beta}^i + P_t^i O_t^i \tag{14}$$

$$L_t^i\left(P_t^i\right) \geq 0 \tag{15}$$

$$P_{min} \leq P_t^i \leq P_{max} \tag{16}$$

$$\boldsymbol{S}^i = \boldsymbol{M} \times \boldsymbol{S}_0^i + \boldsymbol{N}_c \times \mathrm{ReLU}\left(P^i\right)/\boldsymbol{S}_c + \boldsymbol{N}_d \times \mathrm{ReLU}\left(-P^i\right)/\boldsymbol{S}_c \tag{17}$$

$$\boldsymbol{S}_{min} \leq \boldsymbol{S}_t^i \leq \boldsymbol{S}_{max} \tag{18}$$

$$\boldsymbol{S}_T^i = \boldsymbol{S}_{max} \tag{19}$$

$$\boldsymbol{M} = [1 \ 1 \ \dots \ 1]_{1\times 24}^T \tag{20}$$

$$\boldsymbol{N}_c = \begin{bmatrix} \eta & 0 & 0 & \dots & 0 \\ \eta & \eta & 0 & \dots & 0 \\ \eta & \eta & \eta & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \eta & \eta & \eta & \dots & \eta \end{bmatrix}_{24\times 24} \tag{21}$$

$$\boldsymbol{N}_d = \begin{bmatrix} -1/\eta & 0 & 0 & \dots & 0 \\ -1/\eta & -1/\eta & 0 & \dots & 0 \\ -1/\eta & -1/\eta & -1/\eta & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ -1/\eta & -1/\eta & -1/\eta & \dots & -1/\eta \end{bmatrix}_{24\times 24} \tag{22}$$

Eq. (14) represents the calculation of load variable, where $L_{t,\beta}^i$ is the baseline building load, upon which the total load variable L_t^i is derived by adding P_t^i. The first inequality constraint is governed as **Eq. (15)**, ensuring that L_t^i should be greater than or equal to 0. **Eq. (16)** represents another inequality constraint for EV charging and discharging power P_t^i, where P_{min} is −7.4 kW and P_{max} is 7.4 kW. The sequence of SOC $\boldsymbol{S}^i$ is calculated via a state-space model with arrival SOC $\boldsymbol{S}_0^i$, as expressed in **Eq. (17), (20), (21), and (22)**, where η is the battery charging and discharging efficiency, which is set as 0.9 in this study. $\boldsymbol{S}_c$ is the nominal battery capacity, which is assumed as 60 kWh. **Eq. (18)** is the third inequality constraint for EV battery SOC $\boldsymbol{S}^i$. To ensure the longevity of battery, the SOC lower bound $\boldsymbol{S}_{min}$ is 0.2 and upper bound $\boldsymbol{S}_{max}$ is set as 0.9. The formulation integrated real-world characteristics of EV operation, which include EV charging power limits based on empirical analysis of measured smart meter data, battery SOC constraints reflecting battery health considerations, and a departure SOC requirement to ensure next-day travel needs. Additionally, the asymmetric battery charging and discharging efficiency accounts for the typical Li-ion efficiency. These features collectively emulate realistic EV charging behaviors within residential settings. **Eq. (19)** is the hard equality constraint, where we defined the customer leaves home with maximum level of SOC, which is 0.9. It is worth noting that, to simplify the constraint modeling, we assumed that the EV departs at the maximum SOC (0.9) defined. This assumption serves as a consistent boundary condition across all scenarios, ensuring the feasibility of comparisons across models and protecting user driving availability. While this may not reflect the full diversity of real-world user behavior where departure SOC varies with daily travel needs, the proposed framework can readily incorporate user-specific SOC targets when such data is available, such as from mobility surveys or vehicle telemetry. Additionally, the departure SOC could be determined through machine learning algorithms to analyze historical travel patterns and then predict the travel needs, though it was out of the scope of current work. Future work will explore relaxing this assumption and modeling SOC targets as stochastic or scenario-based variables.

Table 1
Marginal CO_2 emission factors.

Hours	1 am	2 am	3 am	4 am	5 am	6 am	7 am	8 am
E_t (kg/kWh)	0.865	0.845	0.81	0.765	0.73	0.70	0.675	0.66
Hours	9 am	10 am	11 am	12 pm	1 pm	2 pm	3 pm	4 pm
E_t (kg/kWh)	0.66	0.68	0.70	0.715	0.72	0.715	0.71	0.705
Hours	5 pm	6 pm	7 pm	8 pm	9 pm	10 pm	11 pm	12 am
E_t (kg/kWh)	0.69	0.685	0.69	0.71	0.74	0.795	0.87	0.87

3.4. DPC framework

In this study, we compared the proposed DPC framework with three alternative methods for handling equality and inequality constraints. The three approaches are: the conventional penalty method, the augmented Lagrangian method, and the KKT method. The conventional penalty method serves as the baseline scenario. Both the augmented Lagrangian method and the KKT method are applied solely to handle equality constraints, with inequality constraints still addressed using the conventional penalty method. Our proposed "Barrier + KKT" method is an enhancement of the KKT method, where the modified barrier function is used to handle inequality constraints, while the KKT network structure ensures the satisfaction of equality constraints. A summary of these methods is provided in Table 2.

The input features were consistent across all four DPC frameworks and included time-related features (time sine and time cosine), baseline building load, occupancy schedule, rate plan, and arrival EV SOC. Normalization was applied to both the baseline building load and the price plan. To facilitate EV discharging from the model outputs, the baseline building load was normalized to the range $[-1,0]$, while the price plan was normalized to $[-1,1]$. The neural network used a batch size of 100.

3.4.1. DPC model 1

The primary distinction between the four DPC frameworks lies in the loss function. For DPC model 1, which employs the conventional penalty method for both equality and inequality constraints, the corresponding loss function is as follows:

$$\ell_1 = \omega_b \ell_b + \omega_c \ell_c + \omega_v \ell_v + \omega_l \ell_l + \omega_p \ell_p + \omega_s \ell_s + \omega_{sf} \ell_{sf} \tag{23}$$

Where ℓ_b, ℓ_c and ℓ_v are the objectives of the problem, which are for minimizing bills, CO_2 and load flattening, respectively. ℓ_l represents the term for load constraint, which is derived from **Eq. (15)**. ℓ_p is the term for EV charging and discharging power constraint, which can be referred to **Eq. (16)**. ℓ_s is the loss from SOC bounds, which can be referred to **Eqs. (17) and (18)**. ℓ_{sf} is the loss term for EV final SOC, which is the equality constraint and can be derived from **Eqs. (17) and (19)**. ω_b, ω_c, ω_v, ω_l, ω_p, ω_s and ω_{sf} are the weights assigned for each of the loss terms.

3.4.2. DPC model 2

As for model 2, the only difference from Model 1 is the treatment of equality constraint. The augmented Lagrangian (AL) method replaces the term ℓ_{sf}. The AL method is a variation of Lagrangian multiplier approach, yet it introduces an additional high-order term to improve convergence. For an optimization problem like **Eq. (6)**, given the equality constraint as:

$$s.t. c_e(x) = 0, \forall i \in \mathcal{E} \tag{24}$$

Where $\boldsymbol{E}$ denotes the indices for equality constraints. The unconstrained problem formulated by AL is then:

$$\min_k \Phi(x) = f(x) + \frac{\upsilon}{2} \sum_e c_e(x)^2 + \sum_e \varphi c_e(x) \tag{25}$$

Where the υ is the AL multiplier for the 2nd order term, and φ is the AL multiplier for the first-order term. Then, for DPC Model 2, the loss function can be expressed as:

$$\ell_2 = \omega_b \ell_b + \omega_c \ell_c + \omega_v \ell_v + \omega_l \ell_l + \omega_p \ell_p + \omega_s \ell_s + AL_{sf} \tag{26}$$

With the multipliers in place, the weight for the equality constraint (final SOC) is no longer necessary. It is important to note that the multipliers are gradually increased through iterations to ensure the satisfaction of the equality constraint during the optimization process. Initial multipliers are 100 and 10 for μ_k and λ_e, respectively.

3.4.3. DPC model 3

For DPC Model 3, since the KKT-projected layers are incorporated into the network, the equality constraint is inherently handled by the network structure and can therefore be excluded from the loss function, where the inequality constraints in DPC Model 3 are still handled using the conventional penalty method:

$$\ell_3 = \omega_b \ell_b + \omega_c \ell_c + \omega_v \ell_v + \omega_l \ell_l + \omega_p \ell_p + \omega_s \ell_s \tag{27}$$

3.4.4. DPC model 4

For DPC Model 4, the equality constraint is still assured by the KKT layer, while the proposed barrier functions are applied to handle the inequality constraints. This integrated approach addresses the limitations of the previous models by providing strong guarantees for both types of constraints. The KKT projection ensures exact satisfaction of the equality constraint, while the barrier functions create powerful gradient signals that guide the optimization away from constraint violations without biasing it toward the interior of the feasible region. By using a shared barrier parameter for all inequality constraints, Model 4 also simplifies the tuning process compared to setting individual penalty weights for each constraint in the conventional approach. The loss function for DPC Model 4 can be denoted as:

$$\ell_4 = \omega_b \ell_b + \omega_c \ell_c + \omega_v \ell_v + \left(\boldsymbol{B}_l + \boldsymbol{B}_p + \boldsymbol{B}_s \right) \tag{28}$$

In this study, the barrier parameter μ is set to be 50. This barrier parameter was selected based on the following criteria: 1) Constraint Enforcement: μ must be sufficiently large to provide steep gradients near constraint boundaries, ensuring violations trigger strong corrective signals; 2) Numerical Stability: Excessively large μ values can cause numerical instability and slow convergence due to extremely steep gradients; 3) Objective Balance: μ should not dominate the objective terms, maintaining meaningful optimization toward the primary goals; and 4) Fair comparison: μ should be the same across the different DPC models to ensure fair comparison. Through empirical testing, we found that 50 achieved the optimal trade-off between constraint satisfaction and objective performance. Setting the same value for the constraint weights assigned for the other DPC models achieved the balance as well. A value below 50 resulted in increased constraint violations, while a

Table 2
Detailed settings for different DPC frameworks.

DPC frameworks		1	2	3	4
Constraints	Equality	Conventional	AL	KKT	KKT
	Inequality	Conventional	Conventional	Conventional	Barrier
Network structure	Hidden layers	3	3	3	3
	Hidden size	64	64	64	64
	Output layers	1	1	2	2
	Activation	Leaky-ReLU	Leaky-ReLU	Leaky-ReLU	Leaky-ReLU
Model training	Optimizer	AdamW	AdamW	AdamW	AdamW
	Learning rate	1e-3	1e-3	1e-3	1e-3
	Max epochs	5000	5000	5000	5000
	Batch size	100	100	100	100
	Early stop	Yes	Yes	Yes	Yes

value above 50 showed dominating effects over the objective performance.

The four approaches are applied to the same problem described previously. It is to be noted that, for fair comparison, all the weights across different models are kept the same. This uniformity in weight allocation ensures that the differences in performance are solely attributable to the distinct methodologies used to handle the equality and inequality constraints, rather than discrepancies in the model configuration or training conditions. Table 3 summarizes the weights assigned to each loss term of each DPC model. The weight parameters in Table 3 were determined through a systematic tuning process aimed at achieving balanced multi-objective performance across all models. The objective weights were selected based on preliminary sensitivity analysis to ensure comparable importance of bill reduction, CO_2 reduction, and load flattening objectives. The constraint penalty weights for Models 1–3 were chosen to provide sufficient penalty strength while avoiding numerical instability. For fair comparison, identical weights were maintained across all models where applicable. Notably, our proposed Barrier+KKT method (Model 4) significantly reduces tuning complexity by consolidating all inequality constraint handling into a single barrier parameter, while eliminating the need for equality constraint weights through the KKT projection layer.

4. Results

This section presents a comprehensive analysis of the optimization results across the four DPC models, comparing their performance in terms of objective achievement and constraint satisfaction. We evaluate the models based on their ability to reduce the electricity bill, decrease CO_2 emission, flatten the load profiles, and minimize the happened times and magnitudes of control violations. Through detailed case studies of individual households and aggregate analyses across all 84 households, we demonstrate the effectiveness of our proposed approach in enhancing both constraint satisfaction and optimization performance.

4.1. Optimization performance for individual households

To provide insight into the behavior of the different DPC models, we first examine the optimization results for two representative households. Fig. 6 and Fig. 7 illustrate the optimization results for one user across all four DPC models, where Fig. 7 serves for a zoomed-in view highlighting constraint violations to aid interpretation of Fig. 6. Fig. 6 provides a visual comparison of how the load is shifted through EV charging and discharging power and how the EV SOC is changed with time under the different DPC models. The top row of the figure shows the load profiles, the middle row displays the EV charging and discharging power, and the bottom row shows the battery SOC trajectories for each DPC model. When comparing the conventional penalty method and the augmented Lagrangian method to the models incorporating the KKT layer, the results reveal that more discharging occurred during on-peak hours in the latter models. Also, the SOC trajectories in the bottom row of Fig. 6 show how the battery charge level evolves throughout the day. All models maintain the SOC within the specified limits, but with varying degrees of compliance. Although the total discharging power appears similar, Models 3 and 4 exhibit less time for EV to discharge compared to Models 1 and 2. The EV was suddenly discharged from 19:00 for Models 3 and 4, instead of discharging continuously during peak hours, as shown for Models 1 and 2. The reason for this scenario is essentially evident in Fig. 7, where it indicates that Models 1 and 2 exhibit significant violations of the load non-negativity constraint, particularly during the peak hours when aggressive discharging is implemented to reduce costs. Models 3 and 4 avoid those violations, which can be directly attributed to the KKT projection layer used for those frameworks. This suggests that the conventional penalty method and augmented Lagrangian approach do not provide sufficient protection against constraint violations when the optimization objectives strongly favor actions that push against those constraints. Also, as seen in the middle row of Fig. 6, Models 1 and 2 resulted in greater discharging power during peak times. In contrast, with the KKT layer applied in Models 3 and 4, discharging is not only concentrated during peak times but also extends to off-peak hours, contributing to a larger reduction in CO_2 emissions. This shift in behavior can be attributed to the reduction in constraint loss terms for Models 3 and 4, which, in turn, increases the weight of the objective loss terms and balances the objectives more justified. This adjustment allows for more flexibility in meeting the objectives, particularly in terms of reducing CO_2 emissions while maintaining the overall system's operational constraints.

Table 3
Weights of each term in different DPC models.

Models	Objectives			Constraints				
	ω_b	ω_c	ω_v	ω_l	ω_p	ω_s	ω_{sf}	μ
Model 1	10	1	5	50	50	50	50	–
Model 2	10	1	5	50	50	50	–	–
Model 3	10	1	5	50	50	50	–	–
Model 4	10	1	5	–	–	–	–	50

The results for a second user, as shown in Fig. 8 and Fig. 9, also follow similar trends but highlight additional aspects of the models' performance. In this case, Fig. 8 shows the similar discharging behavior as the previous user, whereas expanding the majority of discharging periods from peak hours to off-peak hours. Fig. 9 provides a clear view of Fig. 8 where the violations happened, highlighting the superior constraint handling of Model 4 compared to the other approaches. While all models avoid violations of the EV power limits, Models 1, 2, and 3 all exhibit violations of either the load non-negativity constraint or the SOC bounds at various points. From the top row of Fig. 9, it can be observed that Models 1 and 2 violated the load non-negativity. The SOC trajectories for the second household (bottom row of Fig. 9) reveal an important distinction between Models 3 and 4. In Model 3, which uses the KKT layer for equality constraints but conventional penalties for inequality constraints, the SOC temporarily violates the upper limit during the charging period. This violation is visible in Fig. 9, which shows the constraint violations for this household. In contrast, Model 4, which combines the KKT layer with barrier functions, maintains the SOC within the specified bounds throughout the entire period, demonstrating the effectiveness of the barrier function approach in enforcing inequality constraints. Model 4, with its integrated approach to constraint enforcement, shows minimal violations across all constraints, confirming its effectiveness in maintaining feasibility while pursuing the optimization objectives. By providing stronger guarantees of constraint satisfaction, Model 4 allows the optimization to explore more aggressive strategies for meeting the objectives without risking significant constraint violations.

Recent studies on EV charging optimization have shown similar benefits from improved constraint handling mechanisms. For instance, one research demonstrated that robust constraint enforcement is crucial for ensuring the reliability of EV charging controls in the presence of uncertain user behavior [42]. Similarly, another study highlighted the importance of strict constraint satisfaction for maintaining grid stability when integrating large numbers of EVs into distribution networks, as well as the potential to ease the grid's stress by shifting home-charging demand [43]. Our findings align with these studies and extend them by showing that the combination of KKT projection and barrier functions provides a particularly effective approach to constraint management in the DPC framework.

4.2. Aggregated performance analysis

While the individual case studies provide valuable insights into the behavior of the different models, it is also important to analyze their aggregate performance for the sake of load flattening objective. Fig. 10

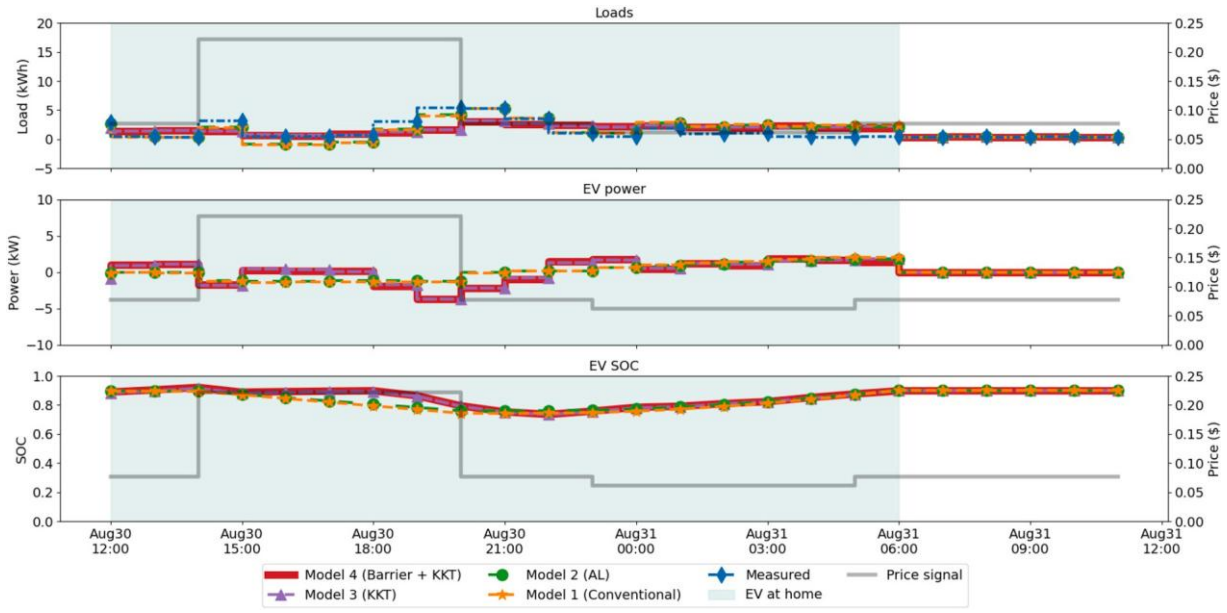

Fig. 6. Optimization results (Example user 1).

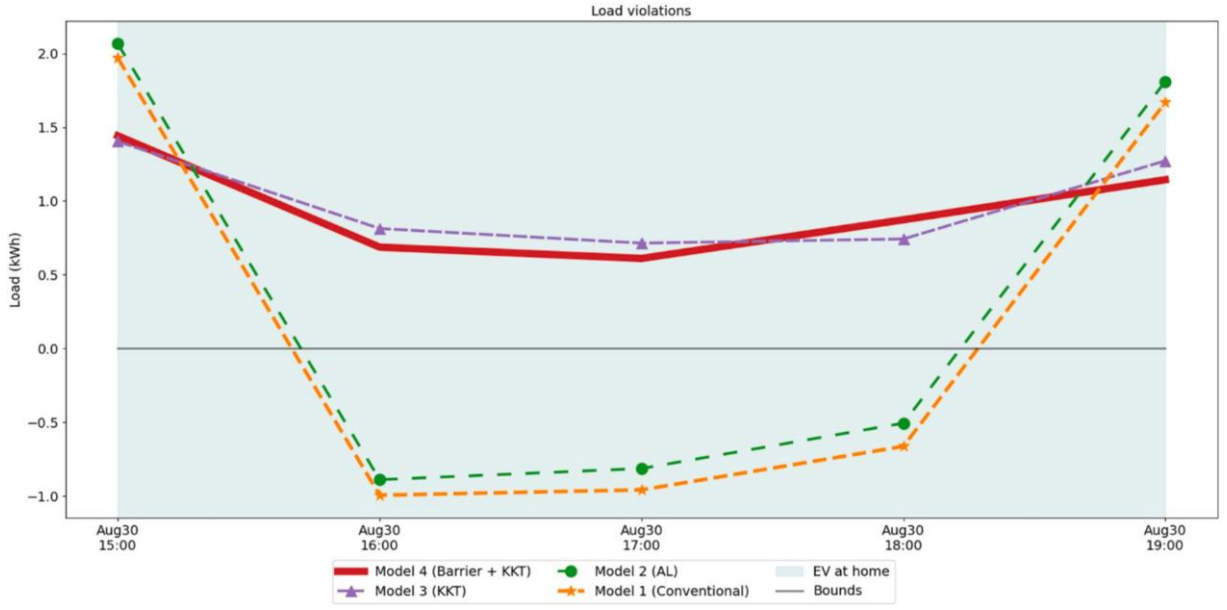

Fig. 7. Constraints violations (Example user 1).

illustrates the aggregated load profile after optimization, highlighting the load flattening objective. The figure compares the original unoptimized load profile (blue) with the optimized profiles generated by each of the four DPC models. It is evident that all four DPC models effectively reduced the peak load, thereby flattening the load pattern. The original load profile shows a significant peak during the evening hours (5 PM to 9 PM), when residential electricity consumption typically reaches its maximum due to the combination of regular household activities and EV charging. After optimization, this peak is substantially reduced across all models, with the load more evenly distributed throughout the day. Among the models, Model 4 achieved the most significant peak load reduction, with a smoother overall profile compared to the other approaches. This improved load flattening can be attributed to the model's ability to better coordinate EV charging and discharging across all households, ensuring that the aggregate load remains more balanced throughout the day. The strategic scheduling of EV charging and discharging behaviors not only reduces costs for individual households but also contributes to overall grid stability by reducing peak demand. The load flattening achieved by our approach has significant implications for grid operations and infrastructure planning. As demonstrated in [44], EV charging coordination can help mitigate the "critical behavior" that emerges when large numbers of EVs charge simultaneously, causing sharp load peaks that stress grid infrastructure. By flattening these peaks, our approach can help reduce the need for grid reinforcements and enable higher EV penetration without requiring costly infrastructure upgrades.

Fig. 11 presents a comprehensive analysis of constraint violations across all 84 households for each model. The boxplots show the distribution of violation magnitudes for three key inequality constraints. From the boxplots, it is clear that the DPC models employing

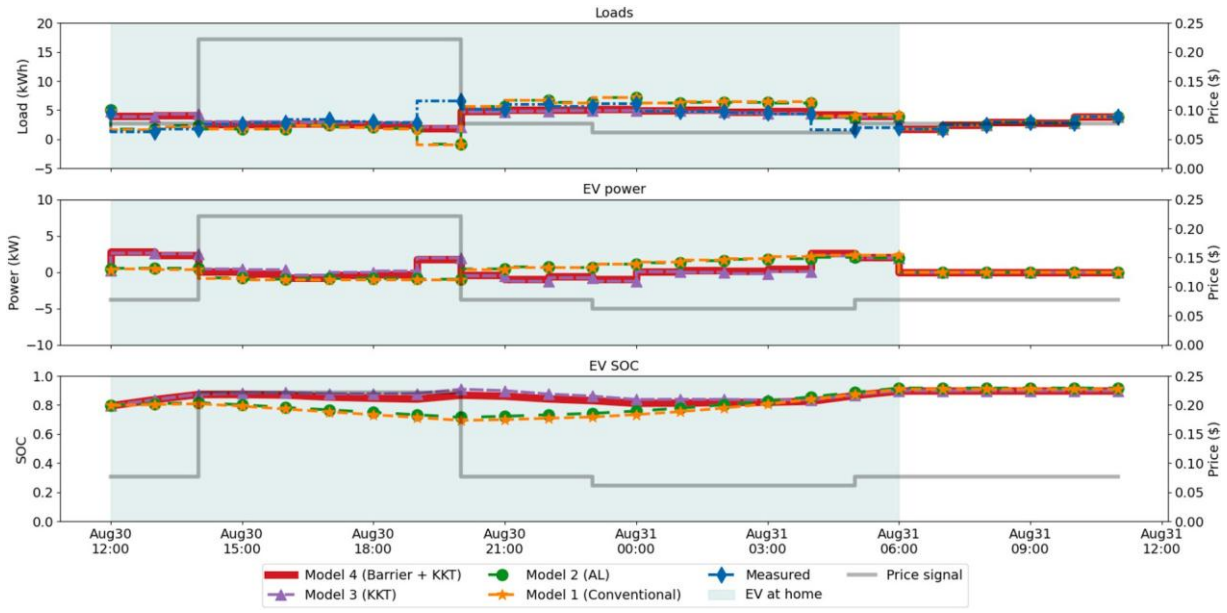

Fig. 8. Optimization results (Example user 2).

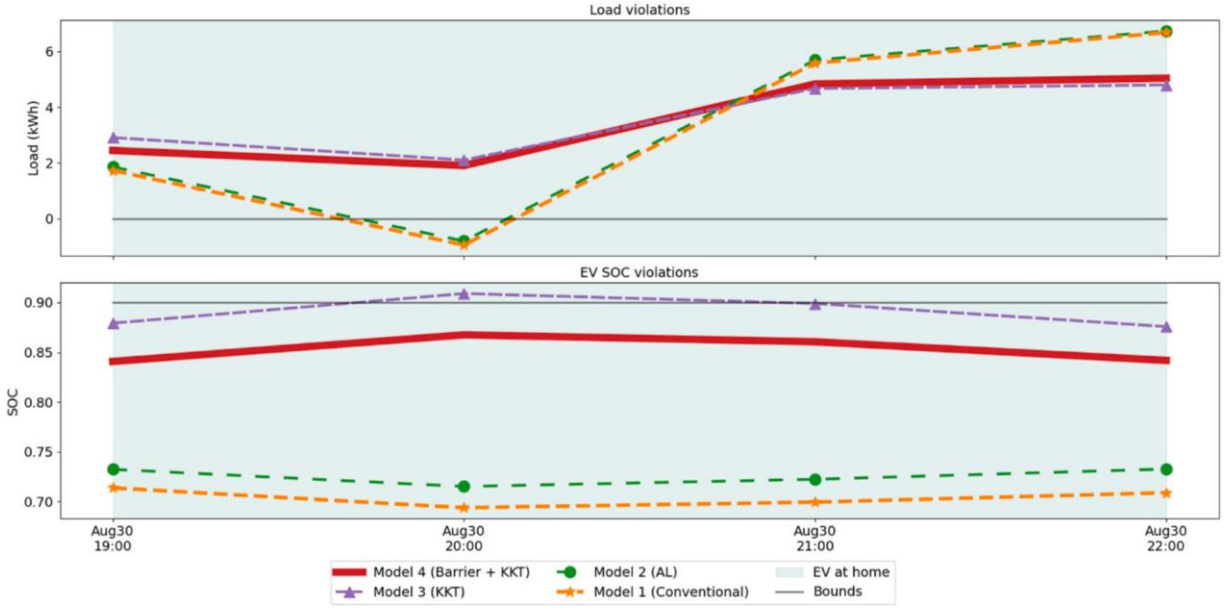

Fig. 9. Constraints violations (Example user 2).

conventional and Lagrangian methods experienced more violations, particularly concerning the inequality constraints for load non-negativity and SOC variables. For the EV power constraint, all models perform well, with no violations observed. This can be attributed to the clear physical limits of EV charging and discharging capabilities, which are easy to enforce even with penalty methods. The load non-negativity constraint shows that Models 1 and 2 are exhibiting frequent violations, while Models 3 and 4 maintain better compliance. However, for the SOC constraint, Models 1 and 2 show significant violations. Model 3 shows fewer violations but with occasionally larger magnitudes, while Model 4 demonstrates excellent performance with minimal violations that are small in magnitude. Overall, the proposed approach (Model 4) resulted in fewer violations, demonstrating superior performance across all constraint categories.

Table 4 provides a quantitative summary of the optimization performance and constraint violation metrics across all four DPC models. Other than the objective metrics, it also presents detailed statistics on constraint violations, including the number of violation instances and the average magnitude of violations for each constraint type. The results in Table 4 show a clear trend of improved performance across all metrics for Models 3 and 4 compared to Models 1 and 2. In terms of all objective metrics, Model 4 achieved the best outcome by highest bill savings, highest CO_2 reductions and highest peak load decreases. The constraint violation statistics provide further evidence of the superior performance of our proposed approach. For the EV power constraint, all models perform well, with no violations found. However, for the other constraint metrics, the differences between models are substantial. Like the SOC constraint, even though Model 3 shows a significant improvement with only 29 violation instances, but with a higher average magnitude. This suggests that while the KKT layer reduced the

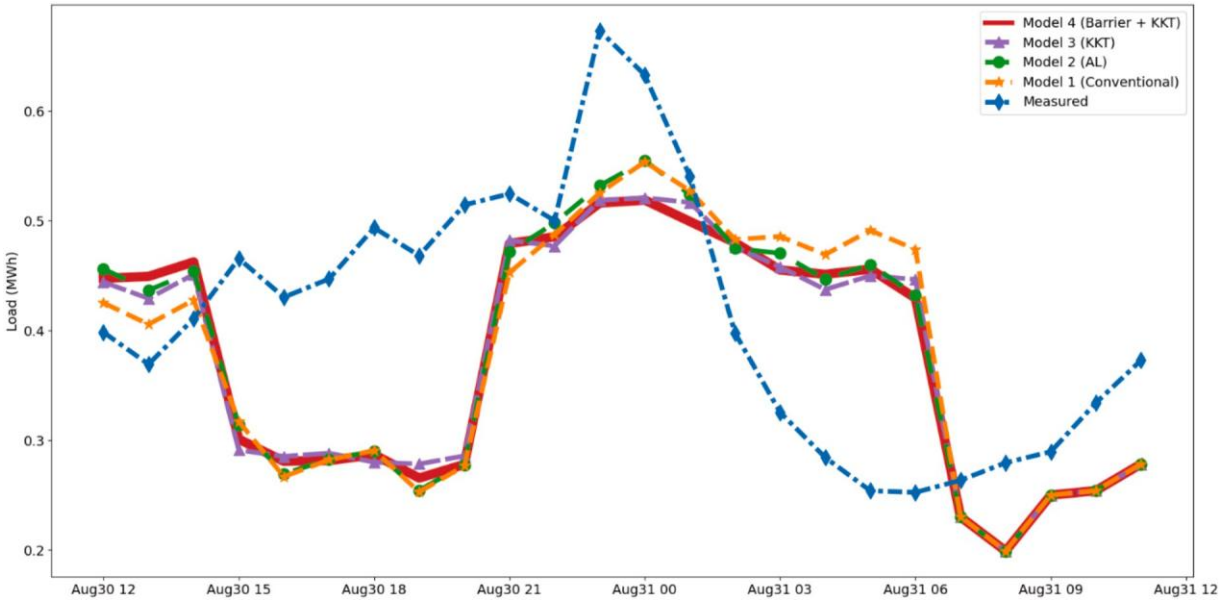

Fig. 10. Aggregated load profiles.

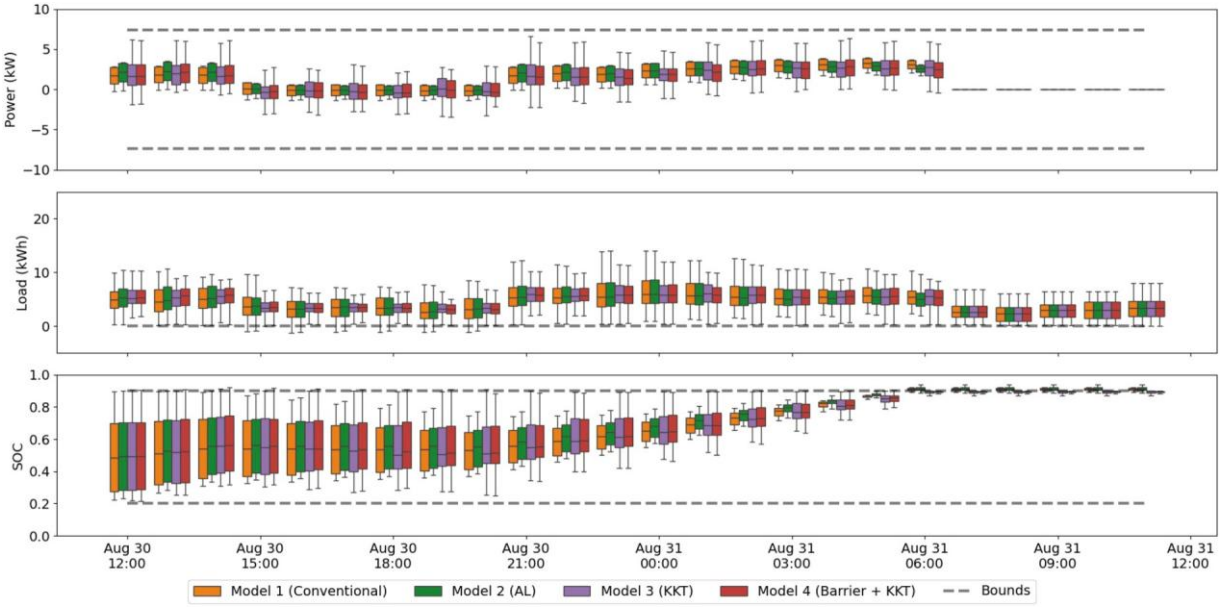

Fig. 11. Boxplot for inequality constraints violations.

Table 4
Overall results for all DPC models.

DPC models			1	2	3	4
Objectives	Bill reductions (%)		16.32	16.63	18.58	19.17
	CO_2 reductions (kg)		556.24	554.50	621.15	625.67
	Peak load reduction (%)		17.75	17.59	22.62	22.94
Violations	EV power	Times	0	0	0	0
		Avg violation (kW)	0	0	0	0
	SOC	Times	468	410	29	6
		Avg violation (%)	1.3	1.7	4.0	0.6
	Load	Times	26	26	0	0
		Avg violation (kWh)	0.87	0.75	0	0
	Final SOC	Avg violation (%)	2.0	1.4	0.8	0.7

effort to account for equality constraint, the conventional penalty method used for inequality constraints sometimes allowed larger violations when they occurred. Model 4, with its combined approach, achieves the best performance in both number of violations and average magnitude. Additionally, for the final SOC, which represents the equality constraint, is effectively handled by the KKT layer. In summary, these quantitative results demonstrate that the proposed approach (Model 4) not only reduces constraint violations but also improves optimization performance across all objectives. By providing stronger guarantees of constraint satisfaction, Model 4 allows the optimization to explore more aggressive strategies for minimizing costs, reducing emissions, and flattening loads, leading to better overall performance compared to the other approaches.

4.3. Computational efficiency

Beyond the optimization performance and constraint satisfaction metrics, we also analyzed the computational efficiency of the four DPC models. During the training process, we applied an early stop strategy to prevent the model from overfitting. Table 5 presents the epochs needed to convergence and total training time for each model. The results show that Models 3 and 4, which incorporate the KKT projection layer, require significantly fewer epochs to converge compared to Models 1 and 2. While the computational cost per epoch is higher for Models 3 and 4 due to the additional projection operations and loss terms, the substantial reduction in the number of epochs required for convergence results in shorter overall training time. This improvement can be attributed to several factors. First, the KKT projection layer provides a direct mechanism for satisfying equality constraint, eliminating the need for the network to learn through iterative penalty adjustments. This allows the optimization to focus primarily on minimizing the objectives, leading to faster convergence. Second, the barrier function approach for inequality constraints provides steeper gradients when approaching constraint boundaries, guiding the optimization more effectively toward feasible solutions compared to conventional penalty methods.

The computational efficiency gains from our proposed approach could be even more significant as the problem scale increases. For larger systems with more households, longer prediction horizons, or more complex constraints, the reduction in training epochs can translate to substantial time savings, making the proposed approach more practical for real-world applications. Additionally, the faster convergence suggests that these approaches are more robust to initialization and hyperparameter settings, requiring less tuning effort to achieve optimal performance.

4.4. Sensitivity analysis

4.4.1. Weekend scenarios

To assess the generalizability of the proposed DPC framework under different user behavior patterns, we conducted an additional simulation using weekend data. This sensitivity study aims to evaluate whether the performance advantages of the proposed DPC formulation (Model 4) hold under conditions that differ significantly from weekday scenarios, specifically in terms of occupancy patterns, charging availability, and energy demand variability. In this simulation, we used the same 84 households' smart meter data on weekends over the month, which included 8 days of meter readings. We used the first 7 days of weekends as training data and last weekend day as the testing set. The occupancy detection and baseline load estimation steps were repeated using the same preprocessing techniques described in Section 3.2. Weekend scenarios typically involve more varied home occupancy, both extremely longer and shorter home presence time windows. Fig. 12 compares the aggregated load profiles across the four models, similar to the weekday analysis. Again, all models demonstrate some peak load flattening, but Model 4 produces the most balanced profile. Fig. 13 presents the boxplots of constraint violations across all households. The results indicate that Models 1 and 2 still suffer from notable violations, particularly for load non-negativity and SOC constraints. The proposed Model 4 maintains strong constraint enforcement.

Table 6 summarizes the performance metrics on the weekend test case. With the most intensive load shifting that happened during peak time for Model 1 results, it achieved the greatest bill reduction. However, this simulation results suffered from notable constraints violations.

Table 5
Computational efficiency comparison.

Model	1	2	3	4
Epochs to convergence	1539	1837	525	541
Total training time (min)	17.7	19.0	11.3	11.0

Model 4 achieved the best CO_2 and peak load reductions, with the best constraint satisfaction without additional judgements. These results further confirmed the robustness of the proposed DPC framework under a different scenario.

4.4.2. Training data size

To evaluate the impact of training data volume on the performance and robustness of the proposed DPC framework, we conducted a sensitivity analysis using two additional training durations for weekday scenarios, which are 7 days and 14 days. For each case, the model was trained from scratch using the specified number of weekday scenarios, while keeping the framework architecture and hyperparameters consistent. All models were evaluated on the same held-out test day. Table 7 detailed the overall performance of Model 4 across the three training sizes. As expected, models trained on fewer days show marginally reduced objective performance and slightly more constraint violations. Despite the different sizes of training dataset, the equality constraint showed similar impact as the average violation for final SOC remained consistent. This sensitivity results highlighted the importance of sufficient training data for achieving high-quality and reliable scheduling outcomes. While the structural constraint-handling design of Model 4 improves robustness, it does not eliminate the dependency on diverse and representative training examples. Using at least 2–3 weeks of training data can ensure stable generalization and safety-compliant performance.

5. Limitations and future work

Our study makes several simplifying assumptions that warrant consideration for practical deployment as the current limitations. Firstly, the assumption of maximum SOC departure represents a conservative upper bound that may not reflect diverse user travel patterns. In addition, the occupancy and arrival time analysis were based on deterministic analysis of smart meter data, neglecting the stochastic nature of real user behavior.

Future work can expand for real implementation. To obtain the user-specific parameters, such as departure SOC, it would require extensive user input, travel pattern learning, or trip planning integration. While our study demonstrates superior constraint handling under deterministic conditions, real-world EV usage involves significant uncertainties including unexpected trips, variable driving patterns, and schedule changes. Comprehensive robustness analysis under such stochastic conditions represents an important area for future investigation that would require dedicated uncertainty modeling, adaptive optimization mechanisms, and validation with real user data. Also, the proposed framework offers a promising architecture for other constrained optimization problems, particularly those involving multi-objective trade-offs and strict feasibility requirements. While this study focuses on EV scheduling, the underlying framework could be extended to domains such as building control, grid dispatch, or industrial energy management, pending further validation in those contexts.

6. Conclusion

This paper proposed a hybrid DPC framework that integrates KKT-based projection layers [35] for equality constraints and modified barrier functions for inequality constraints. The approach enables strict constraint satisfaction while reducing the need for penalty tuning, improving both optimization robustness and computational efficiency. The framework was applied to a real-world EV charging and discharging optimization scenario with multi-objective goals of minimizing electricity bills, reducing CO_2 emissions, and flattening load profiles. Compared to three baseline DPC formulations, the proposed method significantly reduced constraint violations while improving optimization performance, achieving around 3 % greater user's bill reduction, 12 % more CO_2 reduction, and 29 % greater peak load reduction. It also

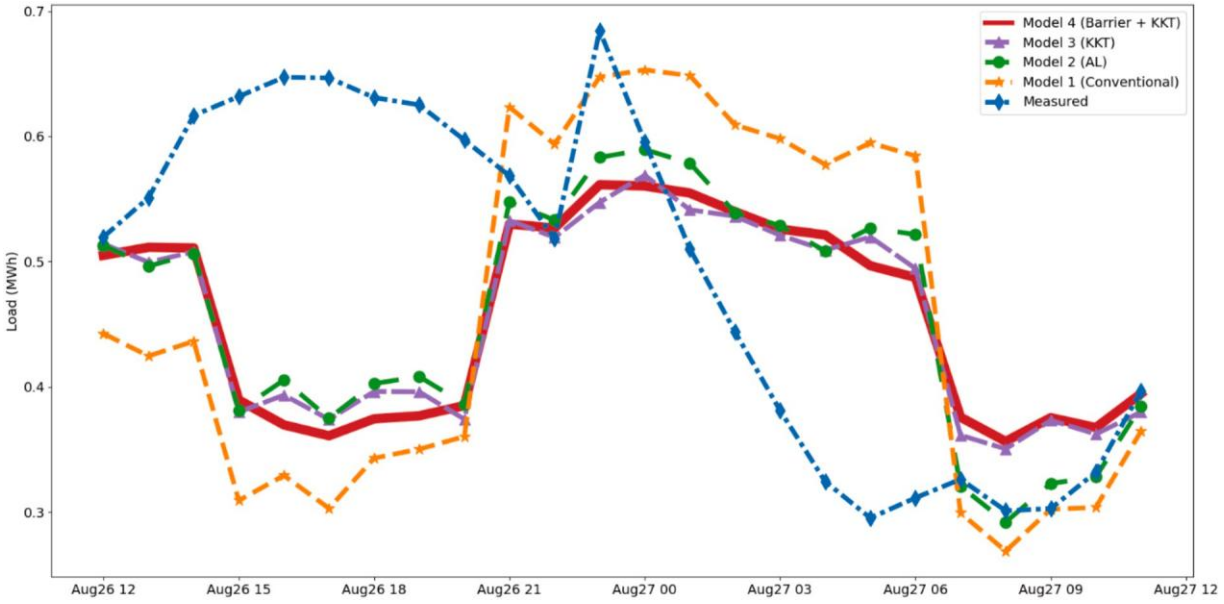

Fig. 12. Aggregated load profiles for weekend scenario.

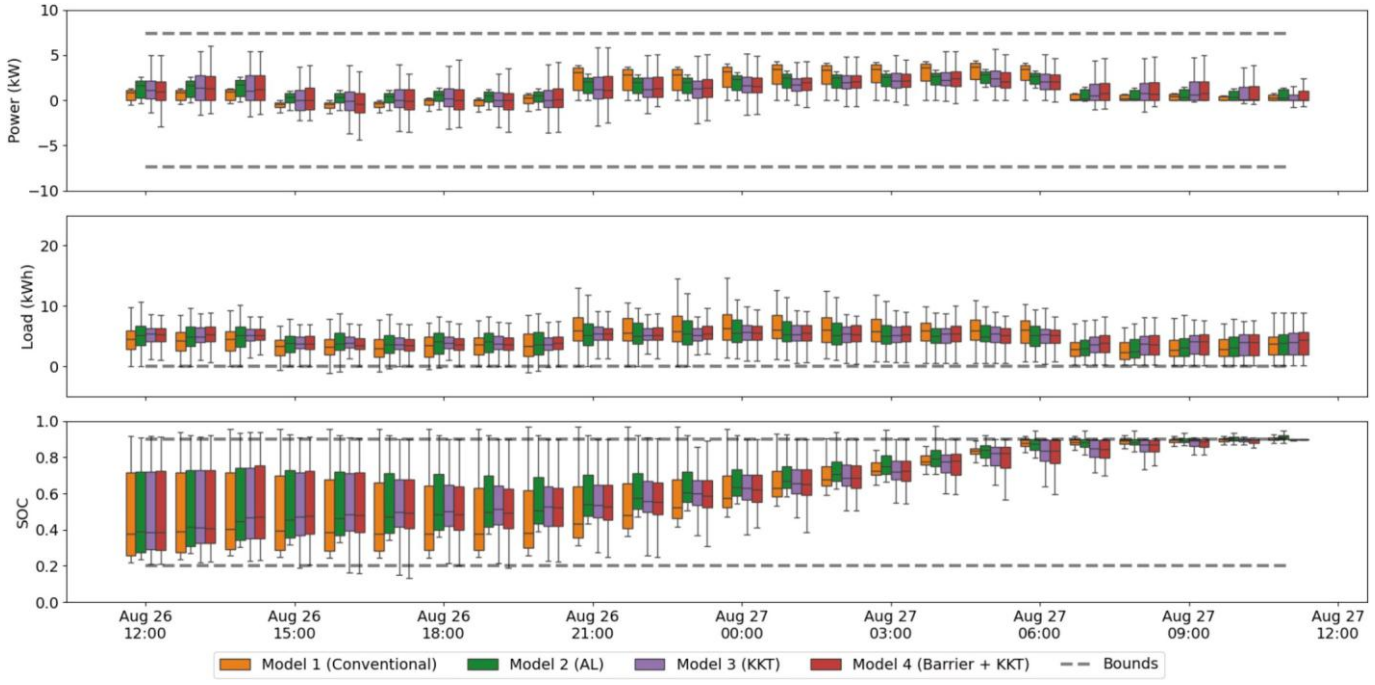

Fig. 13. Boxplot for constraints violations of weekend scenario simulation.

required fewer training epochs, highlighting its data and computational efficiency. In addition, the proposed approach significantly reduces the burden of tuning multiple objective and loss terms, as the equality constraint is naturally embedded in the network structure and the weights for other inequality constraint terms are unified with the barrier parameter. This allows the optimization process to focus on the core objective functions, as the model is better able to direct its resources toward achieving the optimal control performance.

In summary, the Barrier + KKT DPC method not only enhances the model's performance in terms of optimizing EV charging and discharging behaviors but also contributes to a more efficient training process. By minimizing the need for extensive tuning and reducing the time required for model convergence, this approach improves the overall scalability of DPC for solving constrained optimization problems. These combined benefits underscore the potential of the Barrier + KKT DPC framework to streamline the development and application of constrained control policies in various energy systems, offering a practical and efficient solution for complex optimization tasks.

CRediT authorship contribution statement

Yuewei Li: Writing – review & editing, Writing – original draft, Visualization, Methodology, Investigation, Data curation, Conceptualization. **Bing Dong:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Funding acquisition, Conceptualization. **Xuezheng Wang:** Methodology, Conceptualization. **Yueming**

Table 6
Overall results across all DPC models for weekend scenario simulations.

DPC models			1	2	3	4
Objectives	Bill reductions (%)		20.13	16.1	16.63	17.21
	CO_2 reductions (kg)		445.31	503.87	538.51	546.98
	Peak load reduction (%)		4.59	13.88	16.88	17.99
Violations	EV power	Times	0	0	0	0
		Avg violation (kW)	0	0	0	0
	SOC	Times	260	278	39	12
		Avg violation (%)	5.4	5.5	0.3	0.05
	Load	Times	23	7	1	0
		Avg violation (kWh)	0.08	0.03	0.001	0
	Final SOC	Avg violation (%)	1.2	1.8	0.2	0.2

Table 7
Key metrics across 7, 14, and 21-day training size for Model 4.

Training size			7-day	14-day	21-day
Objectives	Bill reductions (%)		12.8	13.52	19.17
	CO_2 reductions (kg)		542.25	614.85	625.67
	Peak load reduction (%)		16.76	21.67	22.94
Violations	EV power	Times	0	0	0
		Avg violation (kW)	0	0	0
	SOC	Times	75	34	6
		Avg violation (%)	1.2	0.6	0.6
	Load	Times	63	59	0
		Avg violation (kWh)	0.44	0.41	0
	Final SOC	Avg violation (%)	0.8	0.8	0.7

Qiu: Writing – review & editing, Supervision, Resources, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Bing Dong reports financial support was provided by National Science Foundation. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This research was supported by the U.S. National Science Foundation (Award No. 2125696).

Data availability

The data that has been used is confidential.

References

[1] Aghemo C, Virgone J, Fracastoro GV, Pellegrino A, Blaso L, Savoyat J, et al. Management and monitoring of public buildings through ICT based systems: control rules for energy saving with lighting and HVAC services. Front Archit Res 2013;2(2):147–61.
[2] Mechri HE, Capozzoli A, Corrado V. USE of the ANOVA approach for sensitive building energy design. Appl Energy 2010;87(10):3073–83.
[3] Drgoňa J, Arroyo J, Figueroa IC, Blum D, Arendt K, Kim D, et al. All you need to know about model predictive control for buildings. Annual Rev Control 2020;50:190–232.
[4] Ma Y, Borrelli F, Hencey B, Coffey B, Bengea S, Haves P. Model predictive control for the operation of building cooling systems. IEEE Trans Control Syst Technol 2011;20(3):796–803.
[5] Oldewurtel F, Parisio A, Jones CN, Gyalistras D, Gwerder M, Stauch V, et al. Use of model predictive control and weather forecasts for energy efficient building climate control. Energ Buildings 2012;45:15–27.
[6] Sturzenegger D, Gyalistras D, Morari M, Smith RS. Model predictive climate control of a swiss office building: implementation, results, and cost–benefit analysis. IEEE Trans Control Syst Technol 2015;24(1):1–12.
[7] Široký J, Oldewurtel F, Cigler J, Privara S. Experimental analysis of model predictive control for an energy efficient building heating system. Appl Energy 2011;88(9):3079–87.
[8] Bianchini, G., Casini, M., Pepe, D., Vicino, A., & Zanvettor, G. G. (2017, December). An integrated MPC approach for demand-response heating and energy storage operation in smart buildings. In 2017 IEEE 56th annual conference on decision and control (CDC) (pp. 3865–3870). IEEE.
[9] Van Cutsem, O., Kayal, M., Blum, D., & Pritoni, M. (2019). Comparison of MPC formulations for building control under commercial time-of-use tariffs (pp. 1–6). IEEE.
[10] Borsche T, Oldewurtel F, Andersson G. Scenario-based MPC for energy schedule compliance with demand response. IFAC Proceedings Volumes 2014;47(3):10299–304.
[11] Esther BP, Kumar KS. A survey on residential demand side management architecture, approaches, optimization models and methods. Renew Sustain Energy Rev 2016;59:342–51.
[12] Fisac JF, Akametalu AK, Zeilinger MN, Kaynama S, Gillula J, Tomlin CJ. A general safety framework for learning-based control in uncertain robotic systems. IEEE Trans Automat Contr 2018;64(7):2737–52.
[13] Asarin E, Bournez O, Dang T, Maler O. Approximate reachability analysis of piecewise-linear dynamical systems. In: International workshop on hybrid systems: Computation and control. Berlin, Heidelberg: Springer Berlin Heidelberg; 2000, March. p. 20–31.
[14] Alessio A, Bemporad A. A survey on explicit model predictive control. Nonlinear Model Predictive Control: Towards New Challenging Applications; 2009. p. 345–69.
[15] Wang X, Dong B. Long-term experimental evaluation and comparison of advanced controls for HVAC systems. Appl Energy 2024;371:123706.
[16] Drgoňa J, Mukherjee S, Tuor A, Halappanavar M, Vrabie D. Learning stochastic parametric diferentiable predictive control policies. IFAC-PapersOnLine 2022;55(25):121–6.
[17] Drgoňa J, Tuor A, Vrabie D. Learning constrained parametric differentiable predictive control policies with guarantees. IEEE Trans Syst Man Cybern Syst 2024;54(6):3596–607.
[18] Mukherjee S, Drgoňa J, Tuor A, Halappanavar M, Vrabie D. Neural lyapunov differentiable predictive control. In: In *2022 IEEE 61st conference on decision and control (CDC)*. IEEE; 2022, December. p. 2097–104.
[19] Drgoňa J, Kiš K, Tuor A, Vrabie D, Klaučo M. Differentiable predictive control: deep learning alternative to explicit model predictive control for unknown nonlinear systems. Journal of Process Control 2022;116:80–92.
[20] Innes, M., Edelman, A., Fischer, K., Rackauckas, C., Saba, E., Shah, V. B., & Tebbutt, W. (2019). A differentiable programming system to bridge machine learning and scientific computing. *arXiv preprint arXiv:1907.07587*.
[21] Chen, B., Cai, Z., & Berg´es, M. (2019, November). Gnu-rl: a precocial reinforcement learning solution for building hvac control using a differentiable mpc policy. In *proceedings of the 6th ACM international conference on systems for energy-efficient buildings, cities, and transportation* (pp. 316-325).
[22] Tumu, R., Cortez, W. S., Drgoňa, J., Vrabie, D. L., & Glavaski, S. (2024). Differentiable predictive control for large-scale urban road networks. *arXiv preprint arXiv:2406.10433*.
[23] Oshin, A., Almubarak, H., & Theodorou, E. A. (2023). Differentiable robust model predictive control. *arXiv preprint arXiv:2308.08426*.
[24] King E, Drgoňa J, Tuor A, Abhyankar S, Bakker C, Bhattacharya A, et al. Koopman-based differentiable predictive control for the dynamics-aware economic dispatch problem. In: In *2022 American control conference (ACC)*. IEEE; 2022, June. p. 2194–201.
[25] Amos B, Jimenez I, Sacks J, Boots B, Kolter JZ. Differentiable mpc for end-to-end planning and control. Adv Neural Inf Processing Syst 2018;31.
[26] Kumar SSP, Tulsyan A, Gopaluni B, Loewen P. A deep learning architecture for predictive control. IFAC-PapersOnLine 2018;51(18):512–7.
[27] Hertneck M, Köhler J, Trimpe S, Allgöwer F. Learning an approximate model predictive controller with guarantees. IEEE Control Syst Lett 2018;2(3):543–8.
[28] Lucia S, Karg B. A deep learning-based approach to robust nonlinear model predictive control. IFAC-PapersOnLine 2018;51(20):511–6.
[29] Wei, T., Zhu, Q., & Maasoumy, M. (2014, November). Co-scheduling of HVAC control, EV charging and battery usage for building energy efficiency. In *2014 IEEE/ACM international conference on computer-aided design (ICCAD)* (pp. 191–196). IEEE.
[30] Wang L, Dubey A, Gebremedhin AH, Srivastava AK, Schulz N. MPC-based decentralized voltage control in power distribution systems with EV and PV coordination. IEEE Trans Smart Grid 2022;13(4):2908–19.
[31] Zakariazadeh A, Jadid S, Siano P. Multi-objective scheduling of electric vehicles in smart distribution system. Energ Conver Manage 2014;79:43–53.
[32] Wan Z, Li H, He H, Prokhorov D. Model-free real-time EV charging scheduling based on deep reinforcement learning. IEEE Trans Smart Grid 2018;10(5):5246–57.
[33] M´arquez-Neila, P., Salzmann, M., & Fua, P. (2017). Imposing hard constraints on deep networks: promises and limitations. *arXiv preprint arXiv:1706.02025*.
[34] Afonso MV, Bioucas-Dias JM, Figueiredo MA. An augmented Lagrangian approach to the constrained optimization formulation of imaging inverse problems. IEEE Trans Image Process 2010;20(3):681–95.
[35] Chen H, Flores GEC, Li C. Physics-informed neural networks with hard linear equality constraints. Comput Chem Eng 2024;189:108764.

[36] Kervadec H, Bouchtiba J, Desrosiers C, Granger E, Dolz J, Ayed IB. Boundary loss for highly unbalanced segmentation. In: International conference on medical imaging with deep learning. PMLR; 2019, May. p. 285–96.
[37] Kervadec H, Dolz J, Tang M, Granger E, Boykov Y, Ayed IB. Constrained-CNN losses for weakly supervised segmentation. Med Image Anal 2019;54:88–99.
[38] Zheng Y, Song Y, Hill DJ, Meng K. Online distributed MPC-based optimal scheduling for EV charging stations in distribution systems. IEEE Trans Industr Inform 2018;15(2):638–49.
[39] McClone G, Ghosh A, Khurram A, Washom B, Kleissl J. Hybrid machine learning forecasting for online mpc of work place electric vehicle charging. IEEE Trans Smart Grid 2023;15(2):1891–901.
[40] Tang W, Zhang YJ. A model predictive control approach for low-complexity electric vehicle charging scheduling: optimality and scalability. IEEE Trans Power Syst 2016;32(2):1050–63.
[41] Chen, D., Barker, S., Subbaswamy, A., Irwin, D., & Shenoy, P. (2013, November). Non-intrusive occupancy monitoring using smart meters. In *proceedings of the 5th ACM workshop on embedded systems for energy-efficient buildings* (pp. 1-8).
[42] Vrettos E, Oldewurtel F, Vasirani M, Andersson G. Centralized and decentralized balance group optimization in electricity markets with demand response. In: In 2013 IEEE Grenoble Conference. IEEE; 2013, June. p. 1–6.
[43] Li Y, Jenn A. Impact of electric vehicle charging demand on power distribution grid congestion. Proc Natl Acad Sci 2024;121(18):e2317599121.
[44] Carvalho R, Buzna L, Gibbens R, Kelly F. Critical behaviour in charging of electric vehicles. New J Phys 2015;17(9):095001.