

MIT Open Access Articles

Learning Model Predictive Control Parameters via Bayesian Optimization for Battery Fast Charging

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Hirt, Sebastian, Höhl, Andreas, Schaeffer, Joachim, Pohlodek, Johannes, Braatz, Richard D et al. 2024. "Learning Model Predictive Control Parameters via Bayesian Optimization for Battery Fast Charging." IFAC-PapersOnLine, 58 (14).

As Published: 10.1016/j.ifacol.2024.08.426

Publisher: Elsevier BV

Persistent URL: <https://hdl.handle.net/1721.1/157696>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution-NonCommercial-NoDerivs



Learning Model Predictive Control Parameters via Bayesian Optimization for Battery Fast Charging

Sebastian Hirt* Andreas Höhl* Joachim Schaeffer**
Johannes Pohlodek* Richard D. Braatz** Rolf Findeisen*

* Technical University of Darmstadt, Darmstadt, Germany

{sebastian.hirt, rolf.findeisen}@iat.tu-darmstadt.de

** Massachusetts Institute of Technology, MA, USA

Abstract: Tuning parameters in model predictive control (MPC) presents significant challenges, particularly when there is a notable discrepancy between the controller's predictions and the actual behavior of the closed-loop plant. This mismatch may stem from factors like substantial model-plant differences, limited prediction horizons that do not cover the entire time of interest, or unforeseen system disturbances. Such mismatches can jeopardize both performance and safety, including constraint satisfaction. Traditional methods address this issue by modifying the finite horizon cost function to better reflect the overall operational cost, learning parts of the prediction model from data, or implementing robust MPC strategies, which might be either computationally intensive or overly cautious. As an alternative, directly optimizing or learning the controller parameters to enhance closed-loop performance has been proposed. We apply Bayesian optimization for efficient learning of unknown model parameters and parameterized constraint backoff terms, aiming to improve closed-loop performance of battery fast charging. This approach establishes a hierarchical control framework where Bayesian optimization directly fine-tunes closed-loop behavior towards a global and long-term objective, while MPC handles lower-level, short-term control tasks. For lithium-ion battery fast charging, we show that the learning approach not only ensures safe operation but also maximizes closed-loop performance. This includes maintaining the battery's operation below its maximum terminal voltage and reducing charging times, all achieved using a standard nominal MPC model with a short horizon and notable initial model-plant mismatch.

Copyright © 2024 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Keywords: Closed-loop Learning, Policy Optimization, Controller Autotuning, Model Predictive Control, Bayesian Optimization, Battery Fast Charging

1. INTRODUCTION

Model predictive control (MPC) allows for the optimal control of linear and nonlinear systems while explicitly taking constraints on the inputs, states, and outputs of the system into account. The closed-loop performance, however, depends on many factors, spanning from the quality of the prediction model used in the controller, appearing disturbances, to a suitable formulation of the cost function and constraints. Selection of these components and of suitable controller parameters is challenging (Lu et al., 2021).

Methods have been proposed to enhance the performance of MPC by exploiting machine learning algorithms (Mesbah et al., 2022; Himmel et al., 2024). Usually, this approach learns a prediction model or parts thereof, e.g., using Gaussian processes (Kocijan, 2016) or neural networks (Höhl et al., 2023). However, a good prediction model does not necessarily lead to good closed-loop performance, depending on the global control objective (Krishnamoorthy, 2023; Sorourifar et al., 2021). Thus, more recently, direct tuning of parametrized MPC policies has been considered. Generally, all components of an MPC

formulation can be parametrized, such as the cost function, the prediction model, and the constraints. Usually, the parameters are tuned either by exploiting reinforcement learning strategies (Zanon and Gros, 2021) or Bayesian optimization (BO) (Paulson et al., 2023; Sorourifar et al., 2021), to improve the closed-loop performance. Typical closed-loop objectives include, e.g., fast convergence to a setpoint or constraint satisfaction. In this work, we consider BO as it is a sample efficient and gradient-free method that is suitable for the optimization of black-box functions and consider it for learning MPC parameters for fast charging of lithium-ion batteries (LIBs) and optimal closed-loop performance.

Learning/tuning of MPC parameters in an outer loop results in a hierarchical framework, where BO offers global optimization of the closed-loop performance over longer periods or multiple episodes while the model predictive controller manages short-term planning. This allows offloading computational resources from the MPC to the BO, where the latter can be run online but does not need to fulfill strict real-time constraints, compared to the MPC. Moreover, the MPC acts as a structured safety layer, ensuring key control theory requirements like closed-loop

stability, robustness, and constraint satisfaction. In this setting, it is crucial to carefully define the global (BO) and local (MPC) cost functions to achieve the desired closed-loop objective.

We demonstrate the approach to fast charging of LIBs, which are essential for storing electrical energy and are a backbone for a carbon-neutral society. Due to high volumetric and gravimetric energy density, LIBs are the storage medium of choice for applications such as consumer electronics, and electric vehicles. However, purchase costs and range anxiety are hurdles for wider adoption of electric vehicles (Egbue and Long, 2012). Optimal fast charging, i.e., charging batteries fast while keeping degradation low (Attia et al., 2020; Matschek et al., 2023), is therefore paramount to avoid oversized battery systems which inflate electric vehicle prices and to allow long trips without excessive charging times.

Much attention has been devoted to applying machine learning to LIBs to address various systems analysis and design problems, e.g., see (Schaeffer et al., 2024b) and citations therein. LIBs are nonlinear dynamical systems, and time-varying due to time- and usage-based degradation, and are an intriguing application for MPC, e.g., (Xavier and Trimboli, 2015; Chen et al., 2021; Lucia et al., 2017; Klein et al., 2011). To keep computational cost reasonable, battery management systems usually use equivalent circuit models (ECMs) for control and estimation of the state-of-charge (SOC) (Plett, 2015).

Here, we apply BO over multiple closed-loop episodes for optimization of a model predictive controller for battery charging. We explore the combination of MPC for ensuring that the battery charging remains within the constraints, with BO to simultaneously learn constraint backups and correct for model-plant mismatch.

The remainder of this article is structured as follows. Section 2 provides an introduction to MPC, BO, and ECM battery models. Subsequently, we present two different case studies in Section 3, followed by the conclusion in Section 4.

2. FUNDAMENTALS

Consider a nonlinear discrete-time dynamical system,

$$\begin{aligned} x_{k+1} &= f(x_k, u_k), \\ y_k &= h(x_k, u_k), \end{aligned} \quad (1)$$

where x_k are the system states, u_k are the system inputs, $f(\cdot)$ are the system dynamics, $h(\cdot)$ is the mapping from states and inputs to the system outputs y_k , and $k \in \mathbb{N}_0$ is the discrete time index.

2.1 Parameterized Model Predictive Control

The objective is to steer system (1) to a desired state (x_d, u_d) while satisfying constraints. One possible approach to do so is MPC, which is based on the repeated solution of a finite-horizon optimal control problem (Rawlings et al., 2017; Findeisen and Allgöwer, 2002). We assume that performance, stability, and repeated feasibility of MPC depends on n_p parameters $\theta \in \Theta \subset \mathbb{R}^{n_p}$, which might influence the cost, the model, or the constraints:

$$\begin{aligned} \min_{\hat{\mathbf{u}}_k} \quad & \left\{ \sum_{i=0}^{N-1} l_\theta(\hat{x}_{i|k}, \hat{u}_{i|k}) + V_{f;\theta}(\hat{x}_{N|k}) \right\} & (2a) \\ \text{s.t.} \quad & \forall i \in \{0, 1, \dots, N-1\}: \\ & \hat{x}_{i+1|k} = \hat{f}_\theta(\hat{x}_{i|k}, \hat{u}_{i|k}), \hat{x}_{0|k} = x_k, & (2b) \\ & \hat{x}_{i|k} \in \mathcal{X}_\theta, \hat{y}_{i|k} \in \mathcal{Y}_\theta, \hat{u}_{i|k} \in \mathcal{U}_\theta, \hat{x}_{N|k} \in \mathcal{E}_\theta, & (2c) \\ & \hat{y}_{i|k} = \hat{h}_\theta(\hat{x}_{i|k}, \hat{u}_{i|k}). & (2d) \end{aligned}$$

Here, $\hat{\cdot}_{i|k}$ denotes the model-based i -step ahead prediction at time index k , $\hat{f}_\theta(\cdot)$ and $\hat{h}_\theta(\cdot)$ define the prediction model, x_k is the measurement of the current system state, N is the length of the prediction horizon, $l_\theta(\cdot)$ and $V_{f;\theta}(\cdot)$ are the stage and terminal cost functions, respectively, and (2c), are the state, output, input, and terminal constraints. Solving (2) yields an optimal input sequence $\hat{\mathbf{u}}_k^* = [\hat{u}_{0|k}^*, \dots, \hat{u}_{N-1|k}^*]$, of which the first element is applied to the system $u_k = \hat{u}_{0|k}^*$. Once new measurements are available, (2) is solved again at the next time index. Similar to (Sorourifar et al., 2021; Lu et al., 2021) we exploit Bayesian optimization to optimize the parameters θ with respect to a closed-loop performance measure.

2.2 Bayesian Optimization

Bayesian optimization (BO) is a sample efficient, global optimization scheme for black-box functions,

$$\theta^* = \arg \max_{\theta \in \Theta} \{G(\theta)\}, \quad (3)$$

where θ^* is the global maximizer of the objective function $G(\cdot)$ on the set Θ . In this work, $G(\cdot)$ is a predetermined but analytically unknown performance measure of the controlled system. The optimization is achieved by sequentially learning a (probabilistic) surrogate model, which is trained on the observed samples of function $G(\cdot)$. As such, a dataset is built sequentially according to $\mathcal{D}_{n+1} = \mathcal{D}_n \cup (\theta_{n+1}, G(\theta_{n+1}))$, $n \in \mathbb{N}_0$, with $\mathcal{D}_0 = (\theta_0, G(\theta_0))$ and the initial parameters θ_0 .

Gaussian processes (GPs) are commonly employed as surrogate models for BO as they tend to be sample efficient (Rasmussen and Williams, 2006). Additionally, GPs not only provide a best estimate (mean) but also an uncertainty measure (variance). The GP is fully defined by its mean function $m(\cdot)$ and covariance function $k(\cdot, \cdot)$ (Rasmussen and Williams, 2006),

$$g(\xi) \sim \mathcal{GP}(m(\xi), k(\xi, \xi')). \quad (4)$$

Assuming a prior mean and covariance function and given the training data \mathcal{D}_n , the posterior mean and variance for the test point θ are given by the inference procedure

$$\begin{aligned} m_n^+(\theta) &= m(\theta) + k(\theta, \vartheta)k_\gamma^{-1}(\gamma - m(\vartheta)), \\ k_n^+(\theta) &= k(\theta, \theta) - k(\theta, \vartheta)k_\gamma^{-1}k(\vartheta, \theta), \end{aligned} \quad (5)$$

with $\vartheta = [\theta_0, \dots, \theta_n]^\top$ and $k_\gamma = k(\vartheta, \vartheta) + \sigma^2 \mathbb{I}$, where σ^2 is the noise variance on the training targets $\gamma = [G(\theta_0), \dots, G(\theta_n)]$. The hyperparameters of the GP are learned by optimizing the logarithmic marginal likelihood (Rasmussen and Williams, 2006).

Based on the surrogate model, an acquisition function $\alpha(\theta)$ is set up, which exploits the mean and covariance of the GP in the search for new parameter samples θ_n . Consequently, in each BO iteration n , the acquisition

function is optimized, yielding the new parameters for the next iteration,

$$\theta_{n+1} = \arg \max_{\theta \in \Theta} \alpha_n(\theta). \quad (6)$$

In this work, we employ the *Expected Improvement* acquisition function (Garnett, 2023).

2.3 Bayesian Optimization for Learning MPC Parameters

For learning/autotuning the parameters of MPC, Bayesian optimization is well suited as it is sample efficient, which reduces the number of necessary closed-loop trials. In this work, we consider tuning of MPC parameters θ for a closed-loop performance objective $G(\theta)$, e.g., distance to a desired state, settling time, or constraint satisfaction. After determining an initial set of parameters θ_0 , a GP surrogate model is built in each optimization iteration n based on the new dataset \mathcal{D}_n . Then the acquisition function is used to specify the next parameter set θ_{n+1} , the closed-loop system is run with parameters θ_{n+1} , and the closed-loop performance $G(\theta_{n+1})$ is evaluated (see Algorithm 1). An advantage of the BO-based MPC tuning is the high sample efficiency and the capability to run the BO online, thus allowing for an adaptive behavior of the MPC to a changing system or environment.

Algorithm 1 MPC Parameter Learning via Bayesian Optimization

Require: Parametrized MPC policy $u_0^*(x; \theta)$, Parameter domain Θ , initial data \mathcal{D}_0 , prior GP mean and covariance functions $m(\cdot)$ and $k(\cdot)$

for $n = 0, 1, 2, \dots$ **do**

Determine posterior GP for \mathcal{D}_n

Maximize acquisition function for θ_{n+1}

Run closed-loop simulation using $u_0^*(x; \theta_{n+1})$

Determine closed-loop performance $G(\theta_{n+1})$

Update data set

end for

2.4 Battery Model and Fast Charging

Lithium-ion batteries (LIBs) are time-varying nonlinear dynamical systems that have been modeled by empirical descriptions (e.g., lookup tables), black-box machine learning (e.g., neural networks), first-principles (e.g., porous electrode theory models), and equivalent circuit models (ECMs) (Schaeffer et al., 2024a). ECMs approximate the battery behavior by an electric circuit using elements such as voltage sources, resistors, capacitors, and more (e.g., Krewer et al. (2018); Schaeffer et al. (2023) and citations therein). An advantage of ECMs over first-principle models is their low computational cost and straightforward implementation, making them suitable for control applications.

The subsequent fast charging case studies use an ECM model with a resistor in series with a resistor capacitance pair (Fig. 1). The voltage source is given by the SOC-dependent cell characteristic open circuit voltage (OCV), R_0 represents the ohmic resistance of the battery, and the R_1, C_1 pair models the polarization of the cell due to the double layer and approximates diffusion.




Fig. 1. R-RC battery ECM model with parameters depending on the SOC z_k .

Consequently, the nonlinear discrete-time dynamics are given by

$$\begin{bmatrix} z_{k+1} \\ U_{1,k+1} \end{bmatrix} = \begin{bmatrix} z_k + \frac{\eta T_s}{Q} I_k \\ (U_{1,k} - R_1 I_k) \exp\left(-\frac{T_s}{R_1 C_1}\right) + R_1 I_k \end{bmatrix}, \quad (7)$$

$$V_{T,k} = \text{OCV} - U_{1,k} - R_0 I_k,$$

where the parameters $R_0(\cdot)$, $R_1(\cdot)$, $C_1(\cdot)$, and the $\text{OCV}(\cdot)$ are functions of the SOC, U_1 is the voltage over C_1 , T_s is the sampling time, η is the Coulombic efficiency, and Q denotes the battery discharge capacity.

We used cubic splines to interpolate the parameters published by Tran et al. (2021), associated with a Samsung SDI INR18650 – 20S lithium nickel manganese cobalt oxide (NMC) cell, and assume ideal cell cooling, i.e., constant temperature, and that the model parameters are independent of the current. As a consequence, the model parameters, i.e., R_0 , R_1 , and C_1 , are only dependent on the SOC z_k , where k is the time index. Batteries change their behavior not only with SOC but also with current and temperature (Hua et al., 2021; Karimi et al., 2023). Consequently, the error between the ECM and the battery increases for operational points away from the parametrization region. In addition, batteries degrade over time and usage, i.e., the capacity and power capabilities decrease. Degradation does not only impact the performance of the application but also control. It furthermore changes the system dynamics, thus leading to a mismatch between model parameters and the cell, introducing a further source of model error.

Fast charging is important to increase the utility of electric vehicles. One approach for designing fast charging protocols is to minimize the charging time subject to degradation constraints, e.g., (Matschek et al., 2023). Commonly, constraints are set for current, voltage, temperature, and lithium-plating overpotential to limit degradation. Here we employ current and voltage constraints given by

$$\forall k \in \mathbb{N}_0 : I_k \leq I_{\max} = 6 \text{ A}, \quad (8)$$

$$\forall k \in \mathbb{N}_0 : 2.5 \text{ V} = V_{T,\min} \leq V_{T,k} \leq V_{T,\max} = 4.2 \text{ V}, \quad (9)$$

with no temperature constraint due to the isothermal assumption. The plating overpotential constraint cannot be set because the chosen ECM does not include plating. Note that we consider current control during the entire charging process due to the model-based approach and do not switch to voltage control as implemented in typical charging circuits, e.g., (Chen and Rincon-Mora, 2006).

Next, we explore how MPC and BO can ensure that the voltage constraint is not violated during fast charging under model-plant mismatch. This article focuses on method-

ology and its demonstration, justifying the simplifications and assumptions mentioned previously.

3. SIMULATION STUDIES

The capabilities of the BO-MPC framework are demonstrated in two case studies focusing on enhancing the closed-loop performance during fast charging while satisfying terminal voltage constraints. In the first case study, the BO optimizes the voltage constraint with a constraint backoff term. In the second case study, the BO optimizes the parameters of the controller's prediction model.

For both case studies, the fast charging problem is formulated as a set point change, where the target state is the fully charged battery, i.e., $z_k = 1$. The MPC is given by

$$\min_{\hat{\mathbf{I}}_k} \left\{ \sum_{i=0}^N (1 - z_{i|k})^2 \right\} \quad (10a)$$

$$\text{s.t. } \forall i \in \{0, \dots, N-1\} : \\ \text{Equation (7)}$$

$$V_{T,\min} \leq \hat{V}_{T,i|k} \leq V_{T,\max}. \quad (10b)$$

We implement the state constraints as soft constraints for numerical reasons. For the sake of simplicity, we employ the ECM model for the controller-intern predictions as well as for the closed-loop simulations. To realize the model-plant mismatch, both case studies use randomly disturbed model parameters in the controller's prediction model. The resulting model parameters differ up to 50% from the true model parameters. The nominal MPC shows a bang-and-ride behavior, i.e., first the current is equal to its upper constraint (bang), but once the system hits the upper voltage constraint, the MPC subsequently continues to “ride” the upper voltage constraint. To avoid constraint violations with a mismatch while fast charging, the global BO objective function is defined as

$$G(\theta) = \sum_{k=0}^M -c_1(1 - z_k)^2 - (\max\{0, V_{T,k} - V_{T,\max}\})^2, \quad (11)$$

where $c_1 \ll 1$ balances the trade-off between minimization of the charging time and constraint satisfaction. Note that the closed-loop values z_k and $V_{T,k}$ in (11) implicitly depend on the parameters θ of the MPC.

3.1 Learning Constraint Backoff

During battery charging, it is essential not to violate the terminal voltage constraint associated with the battery chemistry and given in the battery datasheet. To ensure that the voltage constraints are not violated, a parametrized constraint backoff term is introduced in the MPC via (10b). The resulting output constraint is

$$V_{T,\min} \leq \hat{V}_{T,i|k} \leq V_{T,\max} - b_\theta(z_k), \quad (12)$$

where $b_\theta(\cdot)$ is a backoff function that depends on the current SOC. For efficient implementation, $b_\theta(\cdot)$ is defined by a cubic spline interpolation over a predefined grid. The values at the grid points are stored in $\theta \in \mathbb{R}^7$, where 7 is the chosen number of grid points. The BO adjusts θ to avoid constraint violation while maximizing the state of charge, see (11). The objective is to adjust θ such that the constraint is satisfied after the BO tuning procedure.

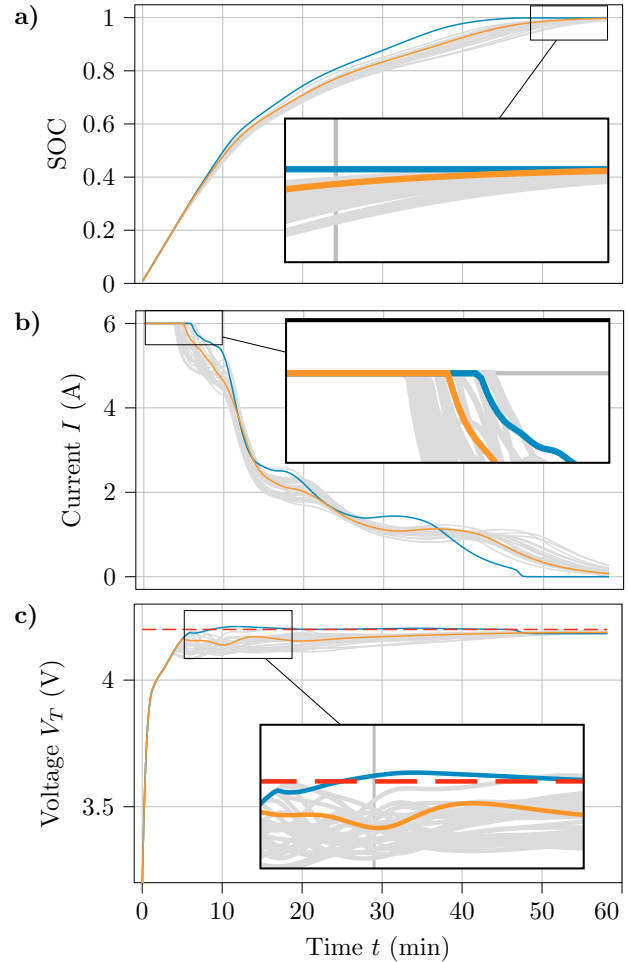


Fig. 2. First case study: Learning constraint backoff. Initial closed-loop solution in blue, BO trials in gray, and optimized closed-loop result in orange. a) SOC trajectories, b) Current trajectories, c) Voltage trajectories with voltage constraint, $V_{T,\max}$, red dashes.

The case study shows that the initial closed-loop solution violates the voltage constraint due to the model-plant mismatch (blue line, Fig. 2c). The final voltage trajectory solution does not violate the voltage constraint because of the learned constraint backoff (orange line, Fig. 2c). During the BO procedure, different parameters for the constraint backoff are sampled in a structured way, where the acquisition function balances exploration and exploitation (gray lines, Fig. 2abc). The SOC trajectory of the optimized closed-loop result rises slower than the initial solution (blue and orange lines in Fig. 2a). While the initial solution promises a faster charge, this is only due to the voltage constraint violation, which results in a later switching time from constant current (CC) to constant voltage (CV). The learned constraint tightening shows that the backoff term decreases with increasing SOC (Fig. 3). The high backoff at the beginning of the charging process ensures that the voltage constraint is not violated at the end of the CC phase. For higher SOC values, however, the constraint backoff is smaller; otherwise, the closed-loop behavior would be overly conservative, and the battery could not be fully charged to SOC = 1. Batteries can only charge if $OCV < V_T$ holds, i.e., the maximum achievable




Fig. 3. Spline interpolation and tuned grid points of the backed-off constraint (green) and OCV (red).

SOC is defined by the intersection of the green and red line in Fig. 3.

As this result is still slightly conservative compared to the initial run, the next case study explores the effect of learning the prediction model parameters.

3.2 Learning Model Parameters

The second case study takes a different approach by learning the parameters of the controller's internal ECM. The BO tunes the values at the grid points of the R_1 spline. To demonstrate that the prediction model does not have to match the plant model exactly for a good closed-loop performance, we only tune the R_1 spline while keeping the mismatch for the other parameters. Note that in general, all of the ECM parameters can be tuned.

The terminal voltage constraint is fulfilled with the tuned system (orange line, Fig. 4c). The initial trial here is identical to the initial trial of the first case study. By tuning the prediction model directly, this case study yields a better closed-loop result, i.e., the maximum SOC is achieved about 10 minutes earlier (Fig. 4a). Additionally, the charging time of the safe tuned result is only slightly lower compared to the unsafe initial result. This is also visible in the tuned current trajectory (orange line, Fig. 4b), which shows a later switching time from CC to CV compared to the initial trajectory. However, the trajectories of the BO trials violate the voltage constraint as there is no explicit upper bound. In contrast, more BO trials keep the voltage constraint in the first case study. Thus, we trade off performance of the final solution with less safe sampled parameter sets.

In summary, the closed-loop behavior significantly improves, despite the fact that the learned model does not converge to the true model (Fig. 5). For SOC ≈ 0.5 , the learned grid points for R_1 (green) are higher than for the initial parameter (purple). This is exactly where the constraint violation occurs at 12 minutes in the initial simulation (Fig. 4). A local backoff is learned through the model, making our results interpretable. For different SOC, the learned grid points might be less intuitive to understand, as they also might be tuned towards compensating the mismatch introduced through the other parameters R_0 and C_1 .

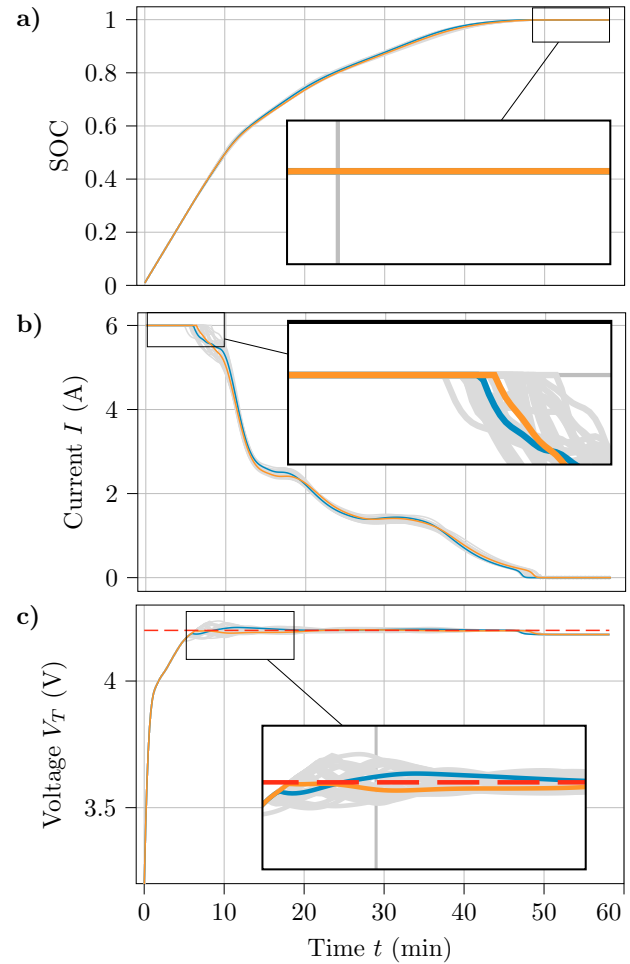


Fig. 4. Second case study: Learning prediction model parameters. Initial closed-loop solution in blue, BO trials in gray, and optimized closed-loop result in orange. a) SOC trajectories, b) Current trajectories, c) Voltage trajectories with voltage constraint, $V_{T,max}$, red dashes.




Fig. 5. Spline interpolation and tuned grid points for R_1 in the prediction ECM. Plant parameter (red), initial prediction parameter (purple), and tuned prediction parameter (green).

4. CONCLUSION

We exploit a hierarchical control framework based on Bayesian optimization for global optimization of a parametrized MPC formulation. While the former takes care of global optimization of the long-term closed-loop behav-

ior, the latter handles lower-level and short-time control tasks. Two specific formulations are proposed within this framework, which are demonstrated in case studies for fast charging of lithium-ion batteries. The first case study learns a constraint backoff term that was parametrized over the SOC of the battery. While mostly safe Bayesian optimization trials were observed, the final result turned out to be conservative compared to the initial unsafe trial. The second case study had less conservative results while showing very similar performance to the initial unsafe trial. While the final result satisfied the constraint, more trials with constraint violations were conducted by the Bayesian optimization in the second case study, indicating more exploration in the parameter space.

REFERENCES

- Attia, P.M., Grover, A., Jin, N., Severson, K.A., Markov, T.M., et al. (2020). Closed-loop optimization of fast-charging protocols for batteries with machine learning. *Nature*, 578(7795), 397–402.
- Chen, G.J., Liu, Y.H., Cheng, Y.S., and Pai, H.Y. (2021). A novel optimal charging algorithm for lithium-ion batteries based on model predictive control. *Energies*, 14(8), 2238.
- Chen, M. and Rincon-Mora, G.A. (2006). Accurate, compact, and power-efficient li-ion battery charger circuit. *IEEE T. Circuits-II*, 53(11), 1180–1184.
- Egbue, O. and Long, S. (2012). Barriers to widespread adoption of electric vehicles: An analysis of consumer attitudes and perceptions. *Energy Policy*, 48, 717–729.
- Findeisen, R. and Allgöwer, F. (2002). An introduction to nonlinear model predictive control. In *21st Benelux Meet. Syst. Control*, volume 11, 119–141. Veldhoven.
- Garnett, R. (2023). *Bayesian Optimization*. Cambridge University Press.
- Himmel, A., Matschek, J., Kok (Louis), R., Morabito, B., Nguyen, H.H., et al. (2024). Chapter 6 - Machine learning for control of (bio)chemical manufacturing systems. In M. Soroush and R. D Braatz (eds.), *Artificial Intelligence in Manufacturing*, 181–240. Academic Press.
- Höhl, A., Maiworm, M., and Findeisen, R. (2023). Path following or tracking model predictive control for towing kites – A question of formulation or learning? In *IEEE Conf. Control Technol. Appl.*, 459–465.
- Hua, X., Zhang, C., and Offer, G. (2021). Finding a better fit for lithium ion batteries: A simple, novel, load dependent, modified equivalent circuit model and parameterization method. *J. Power Sources*, 484, 229117.
- Karimi, D., Behi, H., Van Mierlo, J., and Bercibar, M. (2023). Equivalent circuit model for high-power lithium-ion batteries under high current rates, wide temperature range, and various state of charges. *Batteries*, 9(2), 101.
- Klein, R., Chaturvedi, N.A., Christensen, J., Ahmed, J., Findeisen, R., et al. (2011). Optimal charging strategies in lithium-ion battery. In *Am. Control Conf.*, 382–387. IEEE.
- Kocijan, J. (2016). *Modelling and Control of Dynamic Systems Using Gaussian Process Models*. Springer Cham, Switzerland.
- Krewer, U., Röder, F., Harinath, E., Braatz, R.D., Bedürftig, B., et al. (2018). Dynamic models of li-ion batteries for diagnosis and operation: A review and perspective. *J. Electrochem. Soc.*, 165(16), A3656.
- Krishnamoorthy, D. (2023). On tuning parameterized control policies online for safety-critical systems – Applied to biomedical systems. *IFAC-PapersOnLine*, 56(2), 5781–5786.
- Lu, Q., González, L.D., Kumar, R., and Zavala, V.M. (2021). Bayesian optimization with reference models: A case study in MPC for HVAC central plants. *Comput. Chem. Eng.*, 154, 107491.
- Lucia, S., Torchio, M., Raimondo, D.M., Klein, R., Braatz, R.D., et al. (2017). Towards adaptive health-aware charging of li-ion batteries: A real-time predictive control approach using first-principles models. In *Am. Control Conf.*, 4717–4722. IEEE.
- Matschek, J., Berliner, M.D., Himmel, A., Braatz, R.D., and Findeisen, R. (2023). Necessary optimality conditions for fast lithium-ion battery charging via hybrid simulations. In *Am. Control Conf.*, 3783–3789.
- Mesbah, A., Wabersich, K.P., Schoellig, A.P., Zeilinger, M.N., Lucia, S., et al. (2022). Fusion of machine learning and MPC under uncertainty: What advances are on the horizon? In *Am. Control Conf.*, 342–357.
- Paulson, J.A., Sorourifar, F., and Mesbah, A. (2023). A tutorial on derivative-free policy learning methods for interpretable controller representations. In *Am. Control Conf.*, 1295–1306. IEEE.
- Plett, G.L. (2015). *Battery Management Systems, Volume II: Equivalent-Circuit Methods*. Artech House, Boston.
- Rasmussen, C.E. and Williams, C.K. (2006). *Gaussian Processes for Machine Learning*, volume 2. MIT Press, Cambridge, MA.
- Rawlings, J.B., Mayne, D.Q., and Diehl, M. (2017). *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing Madison, WI.
- Schaeffer, J., Galuppini, G., Rhyu, J., Asinger, P.A., Droop, R., et al. (2024a). Cycle life prediction for lithium-ion batteries: Machine learning and more. In *Am. Control Conf.* Accepted.
- Schaeffer, J., Gasper, P., Garcia-Tamayo, E., Gasper, R., Adachi, M., et al. (2023). Machine learning benchmarks for the classification of equivalent circuit models from electrochemical impedance spectra. *J. Electrochem. Soc.*, 170(6), 060512.
- Schaeffer, J., Lenz, E., Chueh, W.C., Bazant, M.Z., Findeisen, R., et al. (2024b). Interpretation of high-dimensional linear regression: Effects of nullspace and regularization demonstrated on battery data. *Comput. Chem. Eng.*, 180, 108471.
- Sorourifar, F., Makrygirgos, G., Mesbah, A., and Paulson, J.A. (2021). A data-driven automatic tuning method for MPC under uncertainty using constrained bayesian optimization. *IFAC-PapersOnLine*, 54(3), 243–250.
- Tran, M.K., DaCosta, A., Mevawalla, A., Panchal, S., and Fowler, M. (2021). Comparative study of equivalent circuit models performance in four common lithium-ion batteries: LFP, NMC, LMO, NCA. *Batteries*, 7(3), 51.
- Xavier, M.A. and Trimboli, M.S. (2015). Lithium-ion battery cell-level control using constrained model predictive control and equivalent circuit models. *J. Power Sources*, 285, 374–384.
- Zanon, M. and Gros, S. (2021). Safe reinforcement learning using robust MPC. *IEEE Trans. Autom. Control*, 66(8), 3638–3652.