2020 American Control Conference
Denver, CO, USA, July 1-3, 2020

Distributionally Robust Surrogate Optimal Control for Large-Scale Dynamical Systems

Aaron Kandel[1], Saehong Park[1], Hector E. Perez[1], Geumbee Kim[2], Yohwan Choi[2], Hyoung Jun Ahn[2], Won Tae Joe[2], and Scott J. Moura[1]

Abstract— **This paper explores tractable robust optimal control of nonlinear systems with large state spaces. Conventional applications of surrogate modeling for control replace the underlying dynamical model with a data-driven surrogate function. For large-scale systems, this approach possesses a host of shortcomings. We address these challenges by presenting a novel robust surrogate optimization framework for finite-time and receding horizon optimal control. Rather than modeling the entire state transition function, we define a surrogate model which maps the initial state and time series of control inputs to an approximate objective function value. We also define surrogate models which predict time series of relevant constraint functions. Since the bulk of the relevant information is encoded in the initial state, we apply a principal component analysis to project the state onto a reduced basis, allowing surrogate models with tractable parameterizations. To guarantee constraint satisfaction, we use ϕ-divergence to formulate distributionally robust chance constraints which are satisfied for worst-case realizations of the test data modeling error distribution. We validate our approach using a case study of optimal lithium-ion battery fast charging using a large-scale electrochemical battery model.**

I. Introduction

This paper presents a novel algorithmic framework for distributionally robust surrogate optimal control of large-scale dynamical systems. This work is motivated by Bellman's infamous "curse of dimensionality", whereby the computational demands of solving an optimal control problem scale exponentially with the cardinality of the state space of the underlying dynamical system.

In conventional control applications, data-driven surrogate models are often applied to solve nonlinear optimal control problems. These surrogates are typically implemented as replacements or supplements for the underlying state transition models. While this approach presents a host of advantages, it also presents challenges when applied for large-scale dynamical systems. First, modeling state transitions for a dynamical system with hundreds or thousands of state variables can require a surrogate model with an intractable parameterization. Work by Srivastava et al. explores the limitations induced by this shortcoming using data-driven Kriging models [1]. Furthermore, modeling every state transition is often unnecessary and adds significant modeling complexity for limited advantage. This point has been demonstrated in applications of surrogate optimization models, where design variables are mapped directly to an approximate objective function. Modeling errors can also accumulate as the control horizon grows, leading to compounding error and modeling drift that can degrade performance [2]. Thus, for large scale optimal control problems, other methods are frequently pursued. These can include pseudospectral optimal control, control vector parameterization, and approximate dynamic programming (ADP). Each of these methods likewise possess their own unique advantages and shortcomings. For instance, pseudospectral optimal control can require expensive software packages and significant tuning [3]. Control vector parameterization has proven to be a useful tool for large-scale optimal control [4]. However, the inherent approximation implied by simplifying and reducing the control input frequently leads to sub-optimal results. Approximate dynamic programming yields provably optimal results for simple formulations, i.e. linear quadratic problems [5]. However, non-convex optimal control problems can require significant tuning in order to yield useful results. Likewise, control of large-scale systems using ADP can lead to function approximators characterized by computationally intractable parameterizations.

For conventional design optimization problems, surrogate modeling approaches are frequently adopted to map design variables directly to an approximate objective function value. This line of research is particularly popular in the aerospace engineering literature, as aerospace design problems often involve high fidelity and computationally expensive physics-based simulations [6], [7]. This type of approach, despite its utility, fails to completely avoid the underlying curse of dimensionality. Given enough decision variables, the parameterization of the surrogate model can still become untenable for training. With no feasible way to reduce the cardinality of the decision variables, this tool is still only applicable under certain contexts and conditions.

The corollary between this approach and finite-time optimal control would be developing surrogate models which neglect the state transition dynamics entirely. That is, these surrogate models map the initial state and discrete control input sequence directly to an approximate objective function value. However, a significant and unrecognized distinction exists for surrogate optimization models between the applications of control and general design optimization. For control, the bulk of the information is encoded in the initial

This work is funded by LG Chem, and in part supported by a National Science Foundation Graduate Research Fellowship.

[1]Aaron Kandel, Saehong Park, Hector Perez, and Scott Moura are with the Energy, Controls, and Applications Lab (eCal) at the University of California, Berkeley, Berkeley, CA, 94704 USA {aaronkandel, sspark, heperez, smoura}@berkeley.edu

[2]Guembee Kim, Yohwan Choi, Hyoung Jun Ahn, and Won Tae Joe are with LG Chem, BMS Advanced SW Project Team, Yuseong-gu, Daejeon, 305-738, South Korea {gbeekim, nicehwan, ctznahj, wontaejoe}@lgchem.com

state. Consequently, we can use dimensionality reduction techniques like principal component analysis to project the initial state into a reduced basis. This allows us to formulate surrogate optimization models which possess tractable parameterizations, even for particularly large scale dynamical systems. Limited work exists on this approach within the context of control. Peitz et al. provide some brief discussion of surrogate optimal control methods which neglect state transitions, however this discussion is not the focus of their review of multiobjective model-predictive control [8]. Research by Marzat et al. does show some effective nonlinear MPC results using a Kriging surrogate model that neglects state transition dynamics [9]. However, [9] is fundamentally designed and validated for systems with few state variables. Perhaps most importantly, a comprehensive analysis on constraint feasibility is missing from this literature.

We address these open challenges by presenting a novel distributionally robust surrogate optimal control algorithm. First, we define feed-forward neural network surrogate models which map a reduced representation of the initial state and a finite time series of control inputs to an approximate objective function value. A classic approach to enforcing constraints is adding penalty functions to the objective, which can require extensive tuning. Instead, we develop surrogate models which predict time series of the constraint functions based on the same reduced input data. For optimal control problems with a short time horizon, we can obtain approximate solutions by optimizing the model once. However, for optimal control problems on a longer time horizon, we apply these surrogate models within a robust framework. Namely, we reformulate relevant safety constraints as chance constraints which encode statistical distributions of modeling error. To further bolster the robustness of our approach, we apply ϕ-divergence to optimize over the worst case realization of modeling error, within a ball in the space of probability distributions centered around our empirical distribution. Then, the optimal control solution is applied to the plant (i.e., the original full-order model) in the loop to transition into the next state, where we repeat this optimization process until reaching a desired terminal condition. Furthermore, by absorbing the state transition dynamics into the surrogate models, we enable an intuitive application of PCA while also eliminating modeling drift throughout the finite-time optimal control problem.

To validate our approach for a large scale nonlinear dynamical system, we solve the optimal safe-fast charging problem for a lithium-ion battery. Safe lithium-ion battery fast charging is an important application of optimal control, as batteries and charging infrastructure become more powerful. In this paper, we use an electrochemical model known as the single particle model with electrolyte and thermal dynamics (SPMeT). Electrochemical battery models provide more granular information on the state of the battery cell compared to more simple equivalent circuit battery models. This additional information allows us to confidently operate batteries at their safe limits and improve charging times. The cost, however, of using electrochemical models is that they are classic high-dimensional systems characterized by hundreds or even thousands of state variables. This renders optimal control with electrochemisty models a challenge. Some work exists on electrochemical battery model control. Perez et al. apply the Legendre-Gauss-Radau (LGR) pseudo-spectral method with adaptive multi-mesh-interval collocation to solve the optimal fast charging problem using the SPMeT model with a lithium-iron phosphate (LFP) parameterization [3]. Research by Moura et al. and Perez et al. both apply a modified reference governor (MRG) for the fast-charging application. While these results yield a high degree of fast charging performance and safety, optimality is not guaranteed [10], [11]. We demonstrate that our approach yields results similar to those given by [3], [10], but utilizing a more systematic methodology.

II. Problem Formulation

A. Finite-Time Optimal Control Problem Statement

Consider the following nonlinear optimal control problem in discrete time:

$$\min \quad \sum_{k=0}^{N} J_k(x(k), u(k)) \tag{1}$$

$$\text{subject to:} \quad x(k+1) = f(x(k), u(k)) \tag{2}$$

$$g(x(k), u(k)) \leq 0 \tag{3}$$

$$h(x(k), u(k)) = 0 \tag{4}$$

$$x_0 = x(0) \tag{5}$$

where k is the time index of the problem and N is the final time; $x(k) \in \mathbb{R}^n$ is the vector of states at time k; $u(k) \in \mathbb{R}^p$ is the vector of inputs at time k; $J_k(x(k), u(k)) : \mathbb{R}^n \times \mathbb{R}^p \to \mathbb{R}$ is the instantaneous cost at time k as a function of the states and inputs; $f(x(k), u(k)) : \mathbb{R}^n \times \mathbb{R}^p \to \mathbb{R}^n$ represents the linear or nonlinear system dynamics; $g(x(k), u(k)) : \mathbb{R}^n \times \mathbb{R}^p \to \mathbb{R}^m$ represents linear or nonlinear inequality constraints on the states and inputs; and $h(x(k), u(k)) : \mathbb{R}^n \times \mathbb{R}^p \to \mathbb{R}^\ell$ represents linear or nonlinear equality constraints on the states or inputs. We are especially interested in cases where the dimension of x is high, i.e. $n > 10^2, 10^3$, or more, as this confronts us with Bellman's curse of dimensionality.

B. Surrogate Modeling Approach

Conventional applications of surrogate modeling within the controls literature frequently entail replacing the underlying dynamics with a learned surrogate. For large-scale nonlinear dynamical systems, this approach possesses several challenges which we discussed in our introductory section.

This paper adopts a different approach which addresses many of these challenges. First, we replace (1)-(5) with the following optimal control problem statement:

$$\min \quad \mathcal{F}(x_0, U) \tag{6}$$

$$\text{subject to:} \quad \max(\mathcal{G}_i(x_0, U)) \leq 0 \, \forall \, i = 1, \cdots, m \tag{7}$$

Here, the surrogate function $\mathcal{F}$ maps the initial state x_0 and time series of control inputs $U = [u(0), \cdots, u(N)]$ directly to an approximation of the objective function given in (1).

In set notation $\mathcal{F}(\cdot,\cdot) : \mathbb{R}^n \times \mathbb{R}^{p\times(N+1)} \to \mathbb{R}$. Likewise, the surrogate constraint functions $\mathcal{G}_i : \mathbb{R}^n \times \mathbb{R}^{p\times(N+1)} \to \mathbb{R}^{(N+1)}$ take the same inputs and outputs a time series of the relevant constraint function value for each of $i = 1, ..., m$ inequality constraints. Since the output of $\mathcal{G}_i$ is a time series of $g(x(k), u(k))$, we can simply optimize subject to the maximum element in this time series. With a simple modification, we reformulate (7) into distributionally robust chance constraints which probabilistically guarantee safety with respect to modeling errors. We discuss this approach in the following section of this paper.

For a model predictive control application, the optimal control problem in (6)-(7) becomes:

$$\min \quad \mathcal{F}(x_k, U_{k:k+N}) \tag{8}$$

$$\text{subject to:} \quad \max(\mathcal{G}_i(x_k, U_{k:k+N})) \leq 0 \,\forall\, i = 1, \cdots, m \tag{9}$$

At $k = 0$, the initial state becomes the current state, and the control input time series $U_{k:k+N} = [u(k), \cdots, u(k+N)]$ starts at the current state and evolves over a horizon of N time steps into the future. After solving this reduced optimization program, we apply the first control input to the plant model in the loop, simulating one step forward and then repeating the overall process.

Typically, training data for surrogate optimization models is generated via a host of methods. For instance, one popular method in the literature is Latin hypercube sampling (LHS) [7]. For dynamical systems, we can generate suitable training and testing data via random sampling from the state space, and then simulate with random control input trajectories. Overall, the design of experiments can take many forms and is open for further investigation. For unstable dynamical systems, we can randomly generate a set of initial states and simulate the state evolution with random control input sequences over a finite horizon. Compared to episodic simulations over a long time horizon, this approach will help avoid the states diverging, allowing us to obtain useful training data for our surrogate models.

C. State Space Reduction via Principal Component Analysis

For particularly large scale dynamical systems, we can reduce the dimensionality of the initial states using principal component analysis (PCA). Namely, consider the arguments of $\mathcal{F}$ and $\mathcal{G}_i$, $(x_0, U) \in \mathbb{R}^n \times \mathbb{R}^{p\times(N+1)}$. We are interested in $n > 10^2, 10^3$ whereas $p(N+1) \sim 10^1$.

Suppose we have M training data samples for the state x, represented as matrix $X \in \mathbb{R}^{n\times M}$. Consider a so-called "principal component" which can be expressed as:

$$V = w^T X \tag{10}$$

where $w \in \mathbb{R}^{n\times 1}$ is a vector of weights, $V \in \mathbb{R}^{1\times M}$ is an arbitrary principal component. If we consider X as a random matrix, then we seek to choose w to maximize the variance of V

$$\text{var}(V) = w^T X X^T w \tag{11}$$

We then formulate the following optimization problem while constraining w to have unit length:

$$\max_w \quad w^T X X^T w \tag{12}$$

$$\text{subject to} \quad w^T w = 1 \tag{13}$$

which yields the first principal component. This method can be extended to compute multiple principal components, and project the original data onto a reduced basis that maximizes variance [12], thus reducing $x_0 \in \mathbb{R}^n$ to a vector of dimension q, where $q << n$.

III. Robust Chance Constrained Optimization

A. Surrogate Control with Chance Constraints

A fundamental weakness of the surrogate optimal control approximation in (6)-(7) is the accuracy of functions $\mathcal{F}$ and $\mathcal{G}_i$. In particular, the surrogate modeling error of $\mathcal{G}_i$ must be handled with care to ensure solution feasibility. Through some simple statistical analysis, we can rigorously represent feasibility with respect to the original constraints (3) even in the presence of modeling errors. Consider the following surrogate modeling error, or residual, obtained from the test data set used to train $\mathcal{G}_i$:

$$r_i^{(j)} = g_i^{(j)}(x(k), u(k)) - \hat{g}_i^{(j)}(x(k), u(k)), \quad \forall\, k, \forall\, j \tag{14}$$

where $g_i^{(j)}$ is the true constraint function value from the test data, indexed by $j = 1, .., M$ individual data points, and $\hat{g}_i^{(j)}$ is the estimate of this constraint function we obtain from surrogate model $\mathcal{G}_i$. Since our training data (x_0, U) is randomly generated, we can consider residual $r_i^{(j)}$ as samples of a random variable denoted by R_i. For each constraint function $\mathcal{G}_i$, we compute an empirical cumulative distribution function $\hat{\mathbb{P}}$ for the residuals by treating $r_i^{(j)}$ as samples of random variable R_i. We estimate the CDF of R_i empirically from the testing data. Now, rewriting the constraint $g_i(x(k), u(k)) \leq 0$ gives us:

$$g_i(x(k), u(k)) \leq 0, \qquad \forall\, k \tag{15}$$

$$\hat{g}_i(x(k), u(k)) + R_i \leq 0, \qquad \forall, k, \ \forall\ R_i \tag{16}$$

$$\max\{\mathcal{G}_i(x(k), U_{k:k+N})\} + R_i \leq 0, \qquad \forall\ R_i \tag{17}$$

Since R_i is a random variable with potentially unbouded support, we relax this into a chance constraint:

$$\text{Prob}\{\max\{\mathcal{G}_i(x(k), U_{k:k+N})\} + R_i \leq 0\} \geq 1 - \eta_i \tag{18}$$

where η_i is our risk metric, or the allowable risk in violating the chance constraint. This formulation begs the question: *"Why don't we offset the constraint by the maximum residual value?"* This would be equivalent to satisfying (16) with respect to the worst case value of R_i, or that $\eta = 0$. This approach often creates unnecessary conservatism or even an empty feasible set, since the worst case residual is always considered. A practical compromise is chance constraints, which ensures safety for up to the $(1 - \eta_i)$ quantile of the probability distribution, so rare extreme cases

do not unnecessarily compromise performance. We further manipulate (18) using the inverse CDF as follows

$$\text{Prob}\left\{R_i \leq -\max\{\mathcal{G}_i(x(k), U_{k:k+N})\}\right\} \geq 1-\eta_i \quad (19)$$

$$\hat{\mathbb{P}}_{R_i}\left(-\max\{\mathcal{G}_i(x(k), U_{k:k+N})\}\right) \geq 1-\eta_i \quad (20)$$

$$-\max\{\mathcal{G}_i(x(k), U_{k:k+N})\} \geq \hat{\mathbb{P}}_{R_i}^{-1}(1-\eta_i) \quad (21)$$

Importantly, notice that the true distribution for R_i in (19) is approximated by the empirical distribution $\hat{\mathbb{P}}_{R_i}$ in (20), since our only knowledge of the surrogate model approximation error is via test data set residuals. Now, the formulation in (8)-(9) is robustified to:

$$\min \quad \mathcal{F}(x_k, U_{k:k+N}) \quad (22)$$

subject to:

$$\max\{\mathcal{G}_i(x_k, U_{k:k+N})\} + \hat{\mathbb{P}}_{r_i}^{-1}(1-\eta_i) \leq 0, \ \forall\, i = 1, ..., m \quad (23)$$

The inverted CDF $\hat{\mathbb{P}}_{r_i}^{-1}(1-\eta_i)$ takes the form of a constant which adds conservatism with respect to the constraint boundary, which yields an elegant and systematic "safety factor". This conservatism is directly based on empirical data from training and testing the surrogate model $\mathcal{G}_i$. As the surrogate models become more accurate, corresponding to the density function of R_i becoming sharper around zero, then $\hat{\mathbb{P}}_{r_i}^{-1}(1-\eta_i) \to 0$. This is to say that as our models become more accurate, the robust inequality constraint in (23) tightens. This chance constraint approach allows us to guarantee safe control in a probabilistic sense with surrogate models. In the next section, we discuss a data-driven distributionally robust chance constraint (DCC) formulation which improves our guarantee of safety, given limited training and testing data to construct empirical distribution $\hat{\mathbb{P}}_{R_i}$.

B. Data-Driven Distributionally Robust Chance Constraints

The chance constraints provide probabilistic guarantees on feasibility. However, they are based on empirical distributions of surrogate model residuals. In reality, the true distribution of residuals is unknown. To accommodate this distributional, we reformulate the constraint in (20):

$$\inf_{\mathbb{P}\in\mathcal{D}} \mathbb{P}\{\max\{\mathcal{G}_i(x(k), U_k)\} + R_i \leq 0\} \geq 1-\eta_i \quad (24)$$

where $\mathcal{D}$ is a confidence set of potential probability distributions centered around our empirical distribution. Typically $\mathcal{D}$ can be constructed based on *a priori* information. However, in this application the true distribution is unknown.

Since there exists uncertainty in the quality of our empirical distribution $\hat{\mathbb{P}}_{R_i}$, we can characterize the confidence set $\mathcal{D}$ using a distance measure between probability distributions. Borel's law of large numbers suggests that as the amount of data tends towards infinity, the distance between the empirical distribution and the true distribution $\mathcal{D} \to 0$. In this paper, we adopt the ϕ-divergence metric to model the distance between the empirical distribution of residuals and the true distribution of residuals.

1) Quantifying Uncertainty with ϕ-Divergence: We adopt a data-driven distributionally robust chance constraint formulation from [13] to obtain an equivalent reformulation of (24). First, we define ϕ-divergence D_ϕ as:

$$D_\phi(f^*||\hat{f}) = \int_\Omega \phi\left(\frac{f^*(r_i)}{\hat{f}(r_i)}\right)\hat{f}(R_i)dR_i \quad (25)$$

Here, f^* is the true, unknown probability distribution function (pdf), and $\hat{f}$ is the known empirical pdf of residuals, corresponding to probability distributions $\mathbb{P}^*$ and $\hat{\mathbb{P}}$.

Several varying definitions of the function ϕ exist in the literature, each lending itself for different applications and statistical properties. In all cases, ϕ must be a convex function and $\phi(1) = 0$. In this paper, we adopt the Kullbach-Leibler (KL) ϕ-divergence or relative entropy to represent proximity in the probability space. The definition of this function is:

$$\phi_{KL}(x) = x\log(x) - x + 1 \quad (26)$$

We do not know the true probability distribution $\mathbb{P}^*$, but using ϕ-divergence we can assume the true distribution lies within some range $d \in \mathbb{R}$ of our empirical distribution, where d is a distance-like hyperparameter. Based on this assumption, we define a confidence set as follows:

$$\mathcal{D}_\phi = \{\mathbb{P}^* \in \mathcal{P} : D_\phi(f^*||\hat{f}) \leq d, f = \frac{d\mathbb{P}^*}{dR_i}\} \quad (27)$$

In this representation, $\mathcal{P}$ is the set of all feasible probability distributions. Using this confidence set, we can robustify the expression given by (24) as:

$$\inf_{D_\phi(\mathbb{P}^*||\hat{\mathbb{P}})\leq d} \mathbb{P}\{\max\{\mathcal{G}_i(x(k), U_k)\} + r_i \leq 0\} \geq 1-\eta_i \quad (28)$$

Consider the worst-case scenario, where the true probability distribution is a distance d from our empirical distribution. Our chance constraint (18) can be represented as

$$\text{Prob}\{\max\{\mathcal{G}_i(x(k), U_k)\} + r_i \leq 0\} \geq 1-\eta_{KL;i} \quad (29)$$

Here, $\eta_{KL;i}$ is a perturbed risk criterion defined by the following relationship:

$$\eta_{KL;i} = 1 - \inf_{x\in(0,1)}\left\{\frac{e^{-d}x^{1-\eta_i} - 1}{x-1}\right\} \quad (30)$$

This equivalent reformulation is introduced in work by Jiang et al. [13]. By introducing the ϕ-divergence distance-like hyperparameter d, we further pull-in our risk metric to allow additional conservatism based on our confidence in the empirical residuals distribution. In this case, we can solve for $\eta_{KL;i}$ offline using a simple convex optimization program.

This formulation provides a disciplined method for addressing surrogate modeling error, given a distance-like hyperparameter d. Theoretically, as the number of testing data samples $m_{test} \to \infty$, $d \to 0$. This is to say that, as the amount of data comprising our empirical distribution tends towards infinity, our empirical distribution will approach the true distribution ($\hat{\mathbb{P}} \to \mathbb{P}^*$). Past work has explored mapping m_{test} to d with varying approaches that are typically unique

to the underlying case study. In future work, we plan on exploring this relationship for a robust surrogate optimal control methodology that does not require a user-selected hyperparameter d.

IV. Optimal Safe Fast Charging of a Lithium-Ion Battery using an Electrochemical Model

Next we demonstrate the efficacy of our surrogate optimal control approach using an electrochemical lithium-ion battery model, which is a large-scale dynamical system. The fast charging problem has become progressively more important with the proliferation of electrified transport and intermittent renewable power generators. Furthermore, since fast charging applications take batteries to more extreme operating conditions, it is imperative that the model we use for control purposes provides sufficiently detailed information to ensure safety during demanding circumstances. Reduced-order battery models exist, however the lack of granular electrochemical information afforded by such models leads to control policies which are overly conservative or unsafe. These policies (i.e. constant current-constant voltage, or CCCV) therefore provide lower performance in fast charging scenarios. With detailed knowledge of the cell's electrochemical state, we can improve performance and safety simultaneously. However, the key challenge is that electrochemical models typically contain $\sim 10^2$ or $\sim 10^3$ states.

Lithium-ion battery model fidelity varies significantly. Equivalent circuit models (ECM) are perhaps the simplest class of models. However, while their simplicity enables a diversity of control algorithms, they often fail to adequately represent battery cell dynamics during rapid charging charging conditions. For demanding fast charging applications, our modeling objective must include precise predictions of the electrochemical states within the lithium-ion cell. The state information is immensely important when charging at high current, given that the safety and degradation of lithium-ion battery cells depends most strongly on electrochemistry. For fast charging applications, electrochemical models are more suitable, and several versions exist in the literature. In this paper, we conduct optimal control using the single particle model with electrolyte and thermal dynamics (SPMeT). Past work by Perez et al. solves a similar optimal control problem for SPMeT parameterized for a lithium iron phosphate (LFP) cell using pseudospectral optimal control [3]. One challenge with pseudospectral methods is that tuning the initial input trajectory and mesh size is highly non-trivial. Moreover, it's computational complexity prevents real-time applications. This paper uses the results of [3] as one benchmark for our surrogate optimal charging cycle results.

A. Single Particle Model with Electrolyte & Thermal Dynamics

The SPMeT model is derived from the Doyle-Fuller-Newman (DFN) electrochemical battery model. The DFN model employs a continuum of particles in both the anode and cathode of the cell. The SPMeT uses a simplified representation of solid phase diffusion that employs a single spherical particle in each electrode. The governing equations for SPMeT include linear and quasiliniar partial differential equations and a strongly nonlinear voltage output equation, given by:

$$\frac{\partial c_s^{\pm}}{\partial t}(r,t) = \frac{1}{r^2}\frac{\partial}{\partial r}\left[D_s^{\pm} r^2 \frac{\partial c_s^{\pm}}{\partial r}(r,t)\right], \tag{31}$$

$$\varepsilon_e^j \frac{\partial c_e^j}{\partial t}(x,t) = \frac{\partial}{\partial x}\left[D_e^{\text{eff}}(c_e^j)\frac{\partial c_e^j}{\partial x}(x,t) + \frac{1-t_c^0}{F} i_e^j(x,t)\right], \tag{32}$$

where $t \in \mathbb{R}_+$ represents time. Note that superscript j denotes anode, seperator and cathode, $j \in \{+, \text{sep}, -\}$. The terminal voltage output is governed by a combination of electric overpotential, electrode thermodyanmics, and Butler-Volmer kinetics, yielding:

$$\begin{aligned} V(t) = {} & \frac{RT_{cell}(t)}{\alpha F}\sinh^{-1}\left(\frac{I(t)}{2a^+AL^+\bar{i}_0^+(t)}\right) \\ & - \frac{RT_{cell}(t)}{\alpha F}\sinh^{-1}\left(\frac{-I(t)}{2a^-AL^-\bar{i}_0^-(t)}\right) \\ & + U^+(c_{ss}^+(t)) - U^-(c_{ss}^-(t)) \\ & + \left(\frac{R_f^+}{a^+AL^+} + \frac{R_f^-}{a^-AL^-} + \frac{R_{ce}(T_{avg}(t))}{A}\right) I(t) \\ & - \left(\frac{L^+ + 2L^{sep} + L^-}{2A\bar{\kappa}^{eff}}\right) I(t) \\ & + k_{conc}(t)[ln(c_e(0^+,t)) - ln(c_e(0^-,t))], \end{aligned} \tag{33}$$

where c_{ss} is the solid phase surface concentration, namely $c_{ss}^{\pm}(x,t) = c_s^{\pm}(x, R_s^{\pm}, t)$, $U^{\pm}$ is the open-circuit potential, and $c_{s,\max}^{\pm}$ is the maximum possible concentration in the solid phase. The exchange current density i_0^j and solid-electrolyte surface concentration c_{ss}^j are computed as:

$$i_0^j(c_{ss}^j) = k^j\sqrt{c_e^0 c_{ss}^j(t)(c_{s,\max}^j - c_{ss}^j(t))}, \tag{34}$$

$$c_{ss}^j(t) = c_s^j(R_s^j, t), \quad j \in \{+,-\}. \tag{35}$$

The quasilinearity in the electrolyte diffusion equation (32) is rooted in the concentration dependence of the electrolyte diffusion coefficient, $D_e^{\text{eff}}(c_e^j)$.

The nonlinear temperature dynamics are modeled with a single lumped thermal mass subjected to heat transfer:

$$\frac{dT_{cell}}{dt}(t) = \frac{\dot{Q}(t)}{mC_{p;th}} - \frac{T_{cell(t)} - T_\infty}{mC_{p;th}R_{th}} \tag{36}$$

where T_∞ is the ambient temperature, m is the mass of the cell, $C_{p;th}$ is the thermal specific heat capacity of the cell, R_{th} is the thermal resistance of the cell, and $\dot{Q}(t)$ is the heat added from the charging, which is governed by

$$\dot{Q}(t) = I(t)((U^+(SOC_p) - U^-(SOC_n)) - V(t)) \tag{37}$$

Here, $I(t)$ is the input current (the control input), and $V(t)$ is the voltage determined by (33). Both nonlinear open circuit potential functions in (37) are functions of the bulk state of charge (SOC) in the anode and cathode, respectively. This heat generation term makes the temperature dynamics nonlinear. For more details on the SPMeT equations and notation, refer to [14].

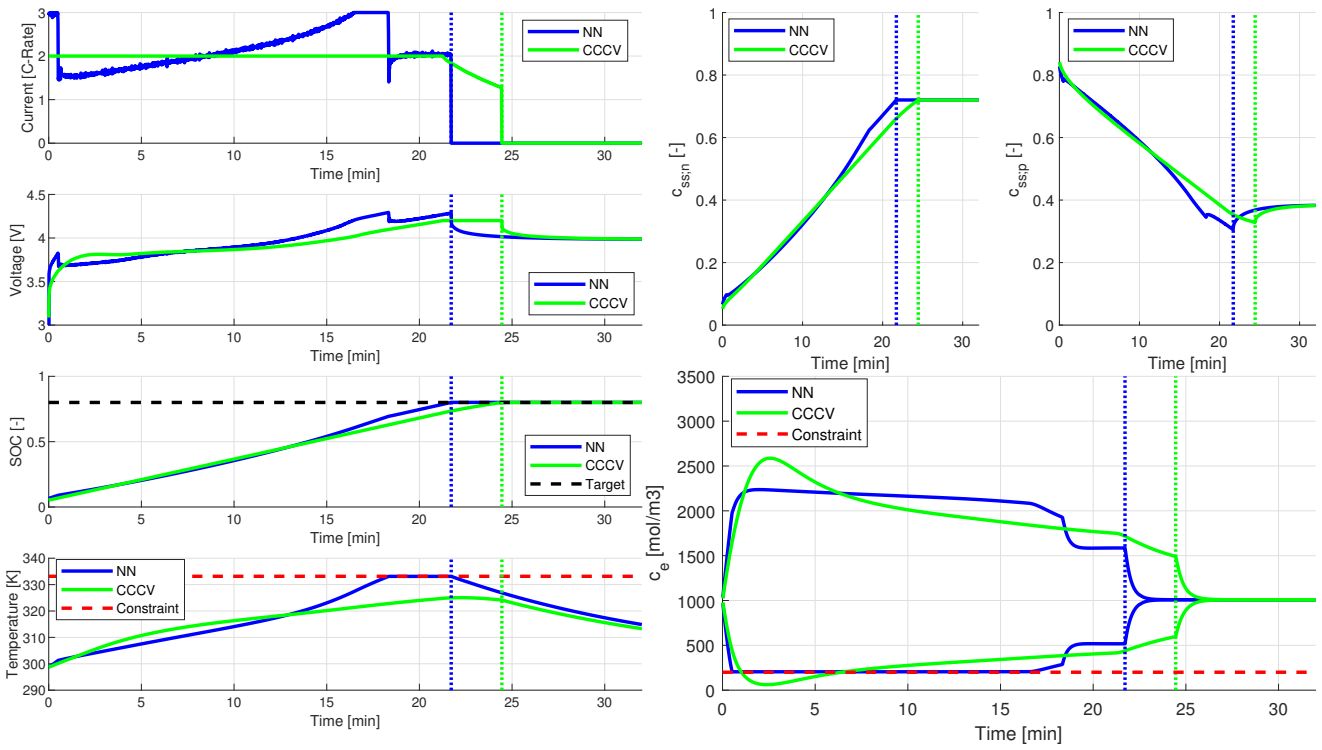

Fig. 1: Optimal charging results for the SPMeT model using a nickel-manganese-cobalt (NMC) cell parameterization. Here, the maximum allowed C-Rate is 3C and the target SOC is 0.8. The initial SOC corresponds to equilibrium conditions at 3.25 Volts.

B. Optimal Control Problem Formulation

In this case study, our objective is to track a target battery state-of-charge:

$$J = \int_{t_0}^{t_F} (SOC_n(t) - SOC_{targ})^2 dt \tag{38}$$

where SOC_n is the normalized bulk concentration in the anode. We define the overall cell SOC as:

$$SOC(t) = \frac{3\int_0^{R_s^-} r^2 c_s^-(r,t)dr}{(R_s^-)^3 c_{s,\max}^- |x_{100\%} - x_{0\%}|} \tag{39}$$

where $|x_{100\%} - x_{0\%}|$ is the stoichiometric difference in the anode. The constraints are:

$$\theta_{\min}^{\pm} \leq \frac{c_s^{\pm}(r,t)}{c_{s,\max}} \leq \theta_{\max}^{\pm} \tag{40}$$

$$c_{e,\min} \leq c_e^l(x,t) \leq c_{e,\max}, l \in -, sep, + \tag{41}$$

$$T_{\min} \leq T_{cell}(t) \leq T_{\max} \tag{42}$$

In this case study, $\theta_{\min}^{\pm} = 0.237$, $\theta_{\max}^{\pm} = 0.863$, $c_{e,\min} = 200$ mol/m^3, $c_{e,\max} = 5000$ mol/m^3, $T_{\min} = 20°$C, and $T_{\max} = 60°$C. The box constraints on the particle surface concentrations, electrolyte concentrations, and cell temperatures are intended to maintain safe charging and prevent rapid cell state-of-health degradation. In this case study, we use individual chance constraints to accommodate these safety requirements. Ideally, we would want to use a joint chance constraint to satisfy all of these constraints with the same probability. However, the inversion of a joint variable CDF function is non-trivial and the subject of on-going research. In future work, we plan on exploring the use of joint chance constraints within our algorithmic architecture. Additionally, we expect to address some minor modeling and parameterization inconsistencies which arose when we identified the parameters of the SPMeT using experimental data. The use of a higher order model like the DFN could potentially resolve this while also revealing new optimal charging protocols.

The PDEs in (31)-(32) are spatially discretized such that we obtain 206 state variables. To generate the training data, we simulate 30 random episodes, each 22 minutes long with random equilibrium initial conditions determined from the equilibrium output voltage. This approach ensures we verify the conservation of lithium principle of the battery model, a property which makes sampling directly from the state space difficult. Each surrogate model is a neural network with 1 hidden layer with 10 neurons with sigmoid activation functions. We reduce the cardinality of the state from 206 to 12 through principal component analysis on the training data, as detailed in Section II-C. PCA reduces the dimensionality of the state used as an input to the surrogate models. We use the finite difference method to discretize the model and conduct moving horizon control where $\Delta t = 1$ second and $N = 4$. Here, $\mathcal{F}$ approximates (38) for a finite, moving

horizon. We construct 6 surrogate constraint models $\mathcal{G}_i$, $i = 1, \cdots, 6$ for each relevant constraint given by (40)-(42). Using the testing data from model training, we obtain empirical modeling error probability distributions which we use to formulate distributionally robust chance constraints. For these constraints, $\eta = 0.02$, $d = 0.001$, and based on a 80-20 training-testing data split $\ell = 7902$. After obtaining an open-loop control policy from our moving horizon control strategy, we apply the first control input to the true SPMeT model (the plant) in the loop, simulate forward one time step and repeat the overall procedure until reaching SOC_{targ}.

C. Optimal Charging Results

Figure 1 shows our results for optimal safe fast charging from an initial SOC of 0.056 to a final SOC of 0.8. In this case, C-rate is a normalized input current rate such that any cell discharged at a constant rate of 1C will be fully discharged from an SOC of 1.0 in exactly 1 hour (assuming no voltage limits or overpotentials). Figure 1 shows optimal charging results with a 3C maximum allowed charging rate for the SPMeT model parameterized for a nickel-manganese-cobalt (NMC) prismatic lithium-ion battery cell. Here, our benchmark is a constant-current constant-voltage (CCCV) charging cycle. The CCCV charging protocol is the most popular fast charging cycle used in industry. It is widely adopted for its simplicity and strong performance. Given CCCV is the industry standard fast charging cycle, it serves as a useful benchmark for our optimal charging results. The NMC cell in our simulations is nominally rated for a maximum CCCV voltage of 4.2 Volts and maximum CCCV current rate of 1C. Within the realm of simulation, we can safely explore the relative performance using a more aggressive 2C CCCV charging rate. We employ a simple $(1 + \lambda)$ evolutionary strategy $(\lambda = 5000)$ in MATLAB to generate the results shown in Figure 1, which required 15 minutes and 5 seconds of computation using a PC equipped with a 9th generation Intel i5 processor with 6 cores. As a result, this control scheme could potentially run in real-time for an actual battery fast charging application.

Figure 1 demonstrates that the optimal charging trajectory based on (38)-(42) follows a CC-Cce-CC-CT profile (constant current, constant anode electrolyte, constant current, constant temperature). This is consistent with prior results obtained in [3], where the optimal solution rides the constraints, and the dominant constraint switches throughout the charging cycle. Compared against the CCCV protocol, there are several noteworthy findings. First, our protocol outperforms the industry standard by charging the battery cell approximately 15% faster than an aggressively tuned CCCV cycle. Perhaps more importantly, by considering the electrochemical states in the charging protocol we ensure internal states stay within a safe operating region. In constrast, CCCV violates $c_{e,\min} \leq c_e^-(x,t)$. Thus, surrogate optimal control will, in principal, extend the useful life of the battery under demanding and aggressive fast charging protocols.

V. Conclusion

This paper presents a novel surrogate modeling approach to optimal control of large scale dynamical systems. We adopt existing ideas from the surrogate optimization and robust optimization literature to develop a novel large-scale optimal control approach. Then, by combining our approach with PCA and distributionally robust optimization we obtain fast surrogate optimal control for large scale systems that is robust to surrogate model uncertainty.We apply our approach to solve the optimal safe-fast charging problem for a lithium-ion battery using the single particle model with electrolyte and temperature dynamics, without reducing the state-space (besides spatial discretization). The added information communicated by the complex electrochemical states allows us to design safe fast charging protocols which exploit electrochemical information to improve overall performance relative to CCCV, the industry standard fast charging protocol.

References

[1] A. Srivastava, K. Hacker, K. Lewis, and T. Simpson, "Development of a kriging based surrogate approximation method for large scale systems," *International Journal for Product and Process Improvement*, 1999.

[2] L. Kaiser, M. Babaeizadeh, P. Mios, B. Osinski, R. H. Campbell, K. Czechowski, D. Erhan, C. Finn, P. Kozakowski, S. Levine, R. Sepassi, G. Tucker, and H. Michalewski, "Model based reinforcement learning for atari," *arXiv*, 2019.

[3] H. Perez, S. Dey, X. Hu, and S. Moura, "Optimal charging of li-ion batteries via a single particle model with electrolyte and thermal dynamics," *Journal of The Electrochemical Society*, vol. 164, no. 7, pp. A1679–A1687, 2017.

[4] M. J. Rothenberger, D. J. Docimo, M. Ghanaatpishe, and H. K. Fathy, "Genetic optimization and experimental validation of a test cycle that maximizes parameter identifiability for a li-ion equivalent-circuit battery model," *Journal of Energy Storage*, vol. 4, pp. 156–166, 2015.

[5] D. P. Bertsekas and J. N. Tsitsiklis, *Neuro-dynamic programming*. Athena Scientific Belmont, MA, 1996, vol. 5.

[6] Y. Mack, T. Goel, W. Shyy, and R. Haftka, "Surrogate model-based optimization framework: A case study in aerospace design," *Evolutionary Computation in Dynamic and Uncertain Systems*, 2007.

[7] N. Queipo, R. Haftka, W. Shyy, T. Goel, R. Vaidyanathan, and P. K. Tucker, "Surrogate-based analysis and optimization," *Progress in Aerospace Sciences*, vol. 41, pp. 1–28, 2005.

[8] S. Peitz and M. Dellnitz, "A survey of recent trends in multiobjective optimal control: Surrogate models, feedback control and objective reduction," *Mathematical and Computational Applications*, vol. 23, no. 30, pp. 1–33, 2018.

[9] J. Marzat and H. Piet-Lahanier, "Design of nonlinear mpc by kriging-based optimization," in *16th IFAC Symposium on System Identification*. Brussels, Belgium: The International Federation of Automatic Control, 2012, pp. 1490–1495.

[10] S. Moura, N. Chaturvedi, and M. Krstic, "Constraint management in li-ion batteries: A modified reference governor approach," in *20133 American Control Conference*. Washington, DC: IFAC, IEEE, 2013.

[11] H. E. Perez, N. Shahmohammadhamedani, and M. S. J. Moura, "Enhanced performance of li-ion batteries via modified reference governors electrochemical model," *EEE/ASME Transactions on Mechatronics*, 2015.

[12] G. C. Calafiore and L. El Ghaoui, *Optimization models*. Cambridge university press, 2014.

[13] R. Jiang and Y. Guan, "Data-driven chance constrained stochastic programs," *Mathematical Programming*, vol. 140, no. 6, pp. 291–327, 2016.

[14] S. J. Moura, F. B. Argomedo, R. Klein, A. Mirtabatabaei, and M. Krstic, "Battery State Estimation for a Single Particle Model with Electrolyte Dynamics," *IEEE Transactions on Control Systems Technology*, vol. 25, no. 2, pp. 453–468, mar 2017. [Online]. Available: http://ieeexplore.ieee.org/document/7489035/