

RESEARCH ARTICLE 

Enhancing dynamical system modeling through interpretable machine-learning augmentations: a case study in cathodic electrophoretic deposition

Christian Jacobsen¹ , Jiayuan Dong², Mehdi Khalloufi² , Xun Huan², Karthik Duraisamy¹ , Maryam Akram³ and Wanjiao Liu³

¹Department of Aerospace Engineering, Ann Arbor, MI, USA

²Department of Mechanical Engineering, Ann Arbor, MI, USA

³Ford Research and Innovation Center, Dearborn, MI, USA

Corresponding author: Christian Jacobsen; Email: csjacobs@umich.edu

Received: 05 June 2024; **Revised:** 07 October 2024; **Accepted:** 11 October 2024


Keywords: e-coat; machine-learning augmentations; neural ODE; uncertainty quantification; variational inference

Abstract

We introduce a comprehensive data-driven framework aimed at enhancing the modeling of physical systems, employing inference techniques and machine-learning enhancements. As a demonstrative application, we pursue the modeling of cathodic electrophoretic deposition, commonly known as e-coating. Our approach illustrates a systematic procedure for enhancing physical models by identifying their limitations through inference on experimental data and introducing adaptable model enhancements to address these shortcomings. We begin by tackling the issue of model parameter identifiability, which reveals aspects of the model that require improvement. To address generalizability, we introduce modifications, which also enhance identifiability. However, these modifications do not fully capture essential experimental behaviors. To overcome this limitation, we incorporate interpretable yet flexible augmentations into the baseline model. These augmentations are parameterized by simple fully-connected neural networks, and we leverage machine-learning tools, particularly neural ordinary differential equations, to learn these augmentations. Our simulations demonstrate that the machine-learning-augmented model more accurately captures observed behaviors and improves predictive accuracy. Nevertheless, we contend that while the model updates offer superior performance and capture the relevant physics, we can reduce off-line computational costs by eliminating certain dynamics without compromising accuracy or interpretability in downstream predictions of quantities of interest, particularly film thickness predictions. The entire process outlined here provides a structured approach to leverage data-driven methods by helping us comprehend the root causes of model inaccuracies and by offering a principled method for enhancing model performance.

Impact Statement

We present an approach to enhancing the modeling of physical systems through a data-driven framework that integrates inference techniques and machine-learning enhancements along with baseline physical models. By focusing on cathodic electrophoretic deposition (e-coating) as a case study, this work systematically identifies limitations in existing physical models using experimental data. The approach presented addresses these

 This research article was awarded Open Materials badge for transparent practices. See the Data Availability Statement for details.

© The Author(s), 2025. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.



limitations by incorporating adaptable model enhancements, significantly improving model parameter identifiability and generalizability. A key contribution of this work is the introduction of interpretable and flexible augmentations to baseline models, parameterized by simple, fully-connected neural networks. Leveraging machine-learning tools, particularly neural ordinary differential equations, the framework effectively learns these augmentations, resulting in simulations that more accurately capture observed behaviors and enhance predictive accuracy. Despite these improvements, the study highlights the potential to reduce off-line computational costs by eliminating certain dynamics without compromising accuracy or interpretability in downstream predictions, particularly for film thickness. This structured approach not only elucidates the root causes of model inaccuracies but also offers a principled method for enhancing model performance.

1. Introduction

Cathodic electrophoretic deposition (EPD), commonly known as e-coating, is a pivotal technique used in industries, such as automotive and manufacturing, to apply protective coatings to various surfaces, preventing corrosion and ensuring durability. Achieving optimal coating properties and process efficiency requires accurate modeling of EPD dynamics, which presents significant challenges due to the complex electrochemical interactions involved.

The primary challenges in modeling the e-coat process are the uncertainty surrounding the physical properties of the coating film during the deposition and a limited understanding of the underlying physics. Various models proposed in the literature have attempted to address these issues (Boyd and Zwack, 1996; Pierce et al., 1981; Rastegar, et al., 2008; Mišković-Stanković, 2002; Ellwood et al., 2009), but they often require empirical calibration based on measurements for film initiation and growth or fail to capture experimental behaviors. In addition, the deposition onset of the coating film, characterized by threshold parameters, remains an open question, leading to discontinuities in model outputs and rendering parameter inference difficult. An extensive study on the modeling of the e-coat process and chemistry involved for constant current and constant voltage can be found in Beck (1976); Pierce, 1981.

E-coating in the automotive industry typically involves employing (Sand, 1899) Eq. (25) to compute the induction time under constant current (CC) conditions. However, the application of a linearly increasing voltage in time is also common in automotive practices to ensure good throw power and prevent defects associated with high voltage (Marlar et al., 2020). To address non-CC and voltage scenarios, an equivalent to Sand's equation is derived using the Laplace transform, assuming a constant diffusion coefficient of the bath solution.

To enhance the modeling of EPD dynamics, specifically in the context of e-coating, this study aims to modify a baseline model to improve model performance with respect to experimentally obtained data. There exist two major types of uncertainty in the baseline model: parametric and model form uncertainties. Parametric uncertainties refer to the uncertainty or lack of precise knowledge about the values of certain parameters in a computational model while model form uncertainty refers to the uncertainty associated with the choice of computational model itself. Parametric uncertainty is often addressed by leveraging experimental data to perform parameter estimation (maximum likelihood estimation, Bayesian inference, etc.), and model form is typically addressed through physical insight and intuition along with data-driven modeling techniques, the latter of which is highlighted in this work. Experimental uncertainties can also significantly affect both the reliability and interpretability of machine-learning methods and models. With increasing levels of noise or error, the learned model parameters may become less accurate, leading to unreliable predictions. However, in this work, we assume that experimental uncertainties are not the primary cause of inaccuracy in our models but rather the model form itself.

Variational inference methods are employed to estimate the parameters of this model in an effort to address the parametric uncertainty associated with the model. However, inherent unidentifiability issues in the model hinder accurate parameter estimation. To overcome these challenges, modifications are introduced to the baseline model to ensure that all parameters are identifiable and more generalizable across experimental conditions, attempting to address model form uncertainty. Nonetheless, the model's inability to capture observed behaviors in experimental data prompts the use of machine-learning

(ML) tools, particularly neural ordinary differential equation (neural ODE) (Chen et al., 2018), to introduce and learn physically relevant and interpretable modifications.

By leveraging neural ODE and carefully crafting learnable augmentations using physical insights, the model's capability to match experimental data is improved while preserving physical interpretability. We use data to enhance an existing full-order dynamical system, effectively creating a surrogate model in the form of model augmentations rather than entirely discovering the physical dynamics (Brunton et al., 2016; Rudy et al., 2017). Our aim is that of improving predictive accuracy while maintaining an interpretable model. As the baseline model is computationally efficient, reduced order modeling techniques (Ahmed et al., 2021; Pawar et al., 2019; San and Maulik, 2018; Hasegawa et al., 2020) are not addressed in this work. To this end, simulations and comparisons with experimental results demonstrate the effectiveness of the improved model in capturing observed behaviors and enhancing the predictive accuracy of the baseline e-coat dynamics model, notably for film thickness prediction, while augmentation forms remain physically interpretable and flexible.

Recent advancements in hybrid machine-learning and data assimilation (DA) techniques have demonstrated the potential to model complex dynamical systems, especially when integrating simulation and observational data. These methods, while showing promise in representing high-dimensional and noisy systems, often prioritize predictive power over interpretability. Our approach aims to maintain a balance between interpretability and accuracy. By leveraging variational inference and neural ODE to design model augmentations, we preserve physical insights while introducing flexibility to model modifications. Hybrid approaches such as those discussed by Cheng et al. (2023). offer complementary strengths, particularly in terms of uncertainty quantification and real-time adaptability. However, our focus on augmenting the physical model with learnable components allows for better adherence to known physical principles, making it suitable for applications where both interpretability and predictive accuracy are critical. This approach aligns with our goal of improving the e-coating model without compromising the transparency of its inner workings. Accurately predicting the polymer film thickness over time is of utmost interest for time and cost savings in industrial applications, and this study argues that not all behaviors present in experimental data are necessary for accurate film thickness predictions. Removing unnecessary complex physical behavior from the model leads to a more efficient formulation of the e-coat process, striking a balance between interpretability, predictive accuracy, and efficiency.

A contribution of this work is to illustrate a principled process of clearly understanding and evaluating the root cause of model shortcomings followed by improving model performance using these insights along with experimental data. Together, our methods can be useful for addressing both parametric and model form uncertainties in computational models where experimental data are available.

The prominent aspects of our approach include the use of variational inference techniques to estimate model parameters from data, an in depth identifiability analysis of parameters in the baseline e-coat model, identifiability and generalizability improvements to the baseline model based on experiment type, and the introduction of learnable ML augmentations to improve model accuracy while retaining interpretability. *Interpretability* in this work refers to the ability of a model to provide insights into its decision-making process in a way that is understandable to humans, specifically by aligning its operations and outputs with known physical laws and principles. This concept goes beyond mere prediction accuracy, aiming to make the internal workings of the model transparent, so that its predictions can be directly related to physical phenomena and understood in terms of physical concepts.

We first illustrate issues with the baseline e-coat dynamics model during inference and prediction which inform the introduction of physically relevant augmentations to improve predictive accuracy. In particular, through variational Bayesian inference, we attempt to infer the parameters of the baseline model and show that identifiability issues hinder the process. Slightly altering the model results in an identifiable model which is more generalizable, but important physically relevant behavior is found to be entirely absent from the model. Finally, we introduce interpretable ML augmentations to improve the accuracy of the baseline model on experimental data. The final model is assessed for both CC and voltage ramp (VR) cases concerning experimental data while trained only on VR cases.

The organization of this manuscript is as follows. Section 2 introduces the baseline model for electro-deposition. Section 3 presents the Bayesian inference methods for estimating the parameters of the baseline model with quantified uncertainty and illustrates the issue of parameter unidentifiability. Section 4 introduces a new model based on a variation of Sand’s equation, valid for any CC/VR experiment, but we show that this model fails to capture all relevant physical behavior observed in experimental data. We then introduce an interpretable ML-based augmentations to the model later in Section 4 to better capture experimental data dynamics and show that removing some of the most complex physical behavior can greatly improve model efficiency without negatively impacting film thickness growth predictions much. Finally, the study concludes in Section 5, highlighting the contributions to e-coating modeling and offering insights into potential applications in diverse fields.

2. Mathematical formulation

In this section, we present the mathematical formulation of the baseline deposition model (Ellwood et al., 2009), a computational framework for solving the baseline model, and the corresponding measurement data acquired from experiments.

2.1. Baseline model

The baseline model is one-dimensional in space, consisting of a cathode and anode placed length L apart and filled with a solution of suspended colloidal particles in between (see Figure 1). After a voltage is applied and prior to film deposition, an electrochemical reaction takes place at the cathode with $2H^+ + 2e^- \rightarrow H_2$ or $2H_2O + 2e^- \rightarrow 2OH^- + H_2$ (Ellwood et al., 2009). As the reaction proceeds, the bath solution basicity increases until a critical pH value is reached and film deposition starts. Film deposition is defined by suspended colloidal particles being deposited on the anode, increasing the circuit resistance.

The film deposition process is separated into three components. First, the electric field within the bath is computed using the conservation of current density given by

$$\nabla \cdot \mathbf{j} = 0 \tag{1}$$

$$\mathbf{j} = -\sigma_{\text{bath}} \nabla \phi \tag{2}$$

$$\phi|_{\Gamma} = R_{\text{film}} j_n \text{ at the interface film-bath,} \tag{3}$$

where \mathbf{j} is the current density, σ_{bath} is the bath conductivity, $j_n = \mathbf{j} \cdot \mathbf{n}$ is the normal component of the current density, ϕ is the electrical potential, R_{film} is the film resistance, and Γ represents the interface between the film and the bath. Second, once film deposition begins, deposition rate is defined as

$$\frac{dh}{dt} = C_v j_n, \tag{4}$$

where h is the film thickness and C_v is the Coulombic efficiency. Third and finally, the film thickness and film resistance are related through

$$\frac{dR_{\text{film}}}{dt} = \rho(j_n) \frac{dh}{dt}, \tag{5}$$




Figure 1. Initial setup for the 1D case.

Table 1. Summary of baseline model parameters

Name	Symbol	Value	Unit
Bath conductivity	σ_{bath}	0.14	S/m
Initial resistance at cathode	R_0	0.5	Ωm^2

where $\rho(\mathbf{j})$ is the film resistivity. In particular, $\rho(\mathbf{j})$ is a decreasing function of the current density, and we adopt an empirical estimate from (Liu and Kevin Ellwood, 2017) of

$$\rho = \max(8 \times 10^5 \exp(-0.1j), 2 \times 10^6). \quad (6)$$

Before deposition begins, the right-hand sides of Eqs. (4) and (5) are both equal to 0. The model defines the deposition onset event according to two criteria: the minimum current density and minimum charge conditions. The onset criteria are critical for accurately predicting the film thickness growth in time. If both of the conditions are met, film deposition begins in the baseline model.

The first condition which determines deposition onset is a minimum current density condition. Once the current density at the cathode reaches a threshold value j_{min} , the film thickness begins to increase. The onset condition parameter j_{min} is unknown and estimated or inferred from experimental data. The second onset condition is a minimum charge condition. The minimum charge criterion assumes that the deposition does not start until the accumulated charge on the cathode reaches a threshold value Q_{min} , with the electric charge Q defined by

$$Q(t) = \int_t j_n dt. \quad (7)$$

The deposition onset threshold Q_{min} is also unknown and estimated or inferred from experimental data. Once both the minimum charge and minimum current conditions are met, film thickness increase as

$$\frac{dh}{dt} = C_v j_n \text{ for } Q > Q_{\text{min}} \text{ and } j > j_{\text{min}}. \quad (8)$$

Additional model parameters have already been estimated from previous experimental data. We thus fix those parameters to values presented in Table 1 and focus only on estimating the key unknown parameters C_v , j_{min} and Q_{min} in the baseline model.

2.2. Computational formulation

We consider two main types of experiments in our computational setup: VR and CC. In VR, the voltage is increased linearly with time at a rate of V_R such that $\phi(t, x=L) = V(t, x=L) = V_R t$. Electric potential is denoted ϕ or V interchangeably in this work. In CC, the current density at the anode is held constant. Experimentally, the current can only be held constant until some maximum voltage V_{max} determined by the available equipment. Numerically, once the maximum voltage is reached in CC, V_{max} is enforced instead of constant current.

Eqs. (1) through (3) are the Poisson equation with Robin boundary condition on the film/bath interface. This model is imposed for both experiments, but the boundary condition at the anode is different for each. Both experiment types are modeled by

$$\sigma_{\text{bath}} \frac{\partial^2 \phi}{\partial x^2} = 0 \text{ in the bath} \quad (9)$$

$$\phi - R_{\text{film}} \sigma_{\text{bath}} \frac{\partial \phi}{\partial x} = 0 \text{ at the interface film-bath,} \quad (10)$$

additionally with VR having an anode boundary condition

$$\phi_{\text{anode}}(t) = \phi_{t=0} + \phi_{\text{ramp}}(t), \tag{11}$$

and CC having an anode boundary condition

$$\sigma_{\text{bath}} \frac{\partial \phi}{\partial x} = j_0. \tag{12}$$

These equations offer an analytic relationship between film resistance and current in 1D for VR:

$$j(t) = \frac{\sigma V(t, x=L)}{\sigma R_{\text{film}}(t) + L}. \tag{13}$$

As $V(t, x=L)$ is set as the boundary condition, we can solve only the resistance dynamics in Eq. (5) and compute $j(t)$ from V and R_{film} . For CC, the current is set by the experimental conditions as j_0 , and the voltage at the anode can be computed by

$$\frac{\partial V}{\partial t} = \frac{\partial R_{\text{film}}}{\partial t} j_0 \text{ for } V < V_{\text{max}}, \tag{14}$$

where V_{max} is the maximum experimental voltage. If $V \geq V_{\text{max}}$ in CC, the relationship between voltage, current, and resistance is given by Eq. (13).

To simulate the baseline dynamics, any black box ODE solver can be implemented to solve Eq. (5) forward in time, computing the time evolution of film resistance. Current, voltage, and charge are given as boundary conditions, computed analytically from Eq. (13), or solved along with resistance dynamics using Eq. (14).

2.3. Experimental setup and data

Experimental data are acquired from a laboratory setup that approximates the computational setting. In the experiment, a 16.0 cm² square anode and cathode are placed at the ends of a long e-coat bath and connected to a power source. A voltage is applied across the anode and the cathode according to either the VR or CC setup. During the course of each experiment, the voltage, current, and film resistance are all measured at a frequency of 10 Hz. For some experiments, thickness measurements are obtained by setting the voltage to zero at some time and measuring the thickness of the film at that time. Note that measuring the thickness terminates the experiment. A depiction of the experimental setup is given in Figure 2

Experiments are performed for both VR and CC at different experimental conditions, for a total of six configurations (see Table 2 for an overview). Multiple trials are repeated for each configuration, where each trial is performed to a different final time in obtaining its thickness measurement. The experimental data are collected in $\mathcal{D} = \{ \{j\}_i, \{R\}_i \}_{i=1}^6$, where $\{j\}_i, \{R\}_i$ respectively represent the sets of current and resistance measurements for each of the six configurations. We do not use voltage data as they can be computed analytically from resistance and current. The data used during inference and learning are truncated for each experiment type i to time t_i such that data have been gathered for at least 3 trials at all times $t \leq t_i$. This is done to provide more accurate estimates of the variance during likelihood




Figure 2. Experimental setup.

Table 2. Descriptions of the experimental data for each of the six experimental conditions

i	Experiment	n_i (# trials)	t_i (s)	Total data points (j, R, V)
1	VR, $V_R = 1.0$ V/s	13	239	31,870
2	VR, $V_R = 0.5$ V/s	13	477	63,620
3	VR, $V_R = 0.125$ V/s	12	639	85,178
4	CC, $j_0 = 10.0$ mA	10	80	10,420
5	CC, $j_0 = 7.5$ mA	10	160	20,820
6	CC, $j_0 = 5.0$ mA	10	240	31,216




Figure 3. Visualization of the $\{j\}_1$ experimental data (all 13 trials) for configuration $VR = 1.0$. Each trial ends at a different time, and data are sampled at a rate of 10 Hz.

computation. An example of current measurements $\{j\}_1$ from the 13 trials under configuration VR, $V_R = 1V$ is shown in Figure 3.

3. Parameter inference using experimental data

Under the standard Bayesian inference framework, the initial knowledge of parameters of interest $\Theta : \omega \rightarrow \mathbb{R}^m$ is described by the prior $p(\theta)$. The observation of data $Y : \Omega \rightarrow \mathbb{R}^n$ given by realizations y are assumed to be related to the parameters through a likelihood $p(y|\theta)$. The posterior $p(\theta|y)$ gives the updated knowledge about parameters θ after observing data y , and is given by Bayes rule:

$$p(\theta|y) = \frac{p(\theta)p(y|\theta)}{p(y)}, \tag{15}$$

where $p(y)$ is known as the evidence and is typically intractable to compute. However, the evidence is constant for some data distribution as does not depend on the parameters θ . In our applications, we are concerned with optimization techniques which require only the log-likelihood and log prior. Thus, the evidence is neglected as a constant and not computed.

In the baseline e-coat model, our parameters of interest are $\theta = \{j_{\min}, C_v, Q_{\min}\}$. Our goal is to infer these parameters given the experimental data described in Section 2.3. The truncated normal distributions are used for the prior to ensure the positive support while still enables the specification of the mode and spread

regarding the parameters. The data random variable Y consists of both current j and film resistance R_{film} . We assume a measurement model of

$$y(\theta, t; \eta) = G(\theta, t; \eta) + \epsilon(t, \eta) \quad (16)$$

where $G(\theta, t; \eta) = \{R_{\text{film}}(\theta, t, \eta), j(\theta, t, \eta)\}$, computed by simulating the baseline model, and is a deterministic function of θ , t , and the experimental configuration parameters η . The measurement noise $\epsilon(t, \eta) \sim \mathcal{N}(0, \Sigma(t, \eta))$ is a function of time and experimental configuration parameters, and we estimate the covariance matrix $\Sigma(t, \eta)$ from experimental data.

Considering our experimental data as \mathcal{D} , the posterior we seek to approximate is given by

$$p(\theta | \mathcal{D}) = \frac{p(\mathcal{D} | \theta)p(\theta)}{p(\mathcal{D})}. \quad (17)$$

The integration involved in directly computing $p(\mathcal{D}) = \int_{\theta} p(\mathcal{D}, \theta) d\theta$ (the “brute force” approach) is expensive to evaluate, particularly in high dimensions. Therefore, numerical techniques have been widely used to approximate the posterior. In addition, computing the posterior directly does not on its own necessarily allow samples to be easily drawn from the posterior. We discuss various inference techniques investigated to perform more efficient inference on the parameters of interest and easily sample from the inferred distributions. The brute force method of directly computing the posterior is referred to as the gridding approach hereafter due to computing the posterior on a discretized grid in the parameter space. Other methods investigated are all forms of variational inference in which the inference problem is transformed to an optimization problem (Duraiamy, 2021; Tran et al., 2021; Guo, 2017; Rezende and Mohamed, 2015).

3.1. Likelihood

Each of the inference methods described require the computation of the (log) likelihood. Given the assumed measurement model in Eq. (16), the likelihood $p(\mathcal{D} | \theta)$ is a Gaussian distribution with mean $G(\theta, t; \eta)$ and covariance matrix $\Sigma(t, \eta)$. We assume that all experiments, trials, and samples within each trial are independent. While the later assumption may not be strictly true, it significantly decreases the computational cost of computing the log likelihood due to a diagonal covariance matrix. The log-likelihood is therefore given by

$$\log p(\mathcal{D} | \theta) = -\frac{1}{2} \sum_{i=1}^6 \sum_{l=1}^{n_i} \sum_{r=1}^{10t_{l,i}} (G(\theta, 0.1r; \eta_i) - \mathcal{D}_{i,l,r})^T \Sigma(0.1r, \eta_i) (G(\theta, 0.1r; \eta_i) - \mathcal{D}_{i,l,r}) + Z, \quad (18)$$

where the time is given by $0.1r$ due to the constant sampling rate of 10 Hz, and $\mathcal{D}_{i,l,r}$ denotes the 2D vector $[\{j\}_i^{(l)}(0.1r), \{R\}_i^{(l)}(0.1r)]^T$, the resistance and current for experiment i in trial l at time $0.1r$. In addition, the matrix $\Sigma(0.1r, \eta_i)$ is diagonal and is computed for experiment i at time $0.1r$ by computing the variance of current and resistance data over all trials which contain data at time $0.1r$. The time $t_{l,i}$ denotes the final experiment time for trial l of experimental configuration i . Finally, Z is the negative log of the likelihood normalization constant which does not depend on θ .

3.2. Gridding approach

Computing the Bayesian posterior is typically prohibitively expensive in practice. For our application, simulating the baseline model is the dominant bottleneck in terms of time complexity. Evaluating the forward model as little as possible will provide the greatest benefit for the efficiency of the inference process. More efficient methods of approximating the posterior are thus investigated to perform parameter inference. Each of these results is compared to the Bayesian posterior using a gridding approach, which is considered the “true” posterior.

We employ a simple gridding approach to compute the Bayesian posterior in which the parameter space is discretized into a d -dimensional grid of uniform spacing. This grid is then sequentially refined with higher resolution near the maximum a posteriori (MAP) and any modes of the Bayesian posterior.

It is often computationally more stable to compute the log-distributions first, rather than the distribution directly. Considering a single grid point θ_i , we first compute the quantity

$$\log p(\theta_i|y) + \log p(y) = \log p(\theta_i) + \log p(y|\theta_i) \quad (19)$$

at all grid points in the parameter space. Note that we ignore the *evidence* term $p(y)$ in the gridding approach until Eq. (19) is computed at each point. It is assumed that the grid bounds in parameter space are sufficiently large to capture the most significant aspects of the distribution.

The final distribution is computed at each point in the grid by normalizing the quantity computed in Eq. (19) using

$$p(\theta_i|y) = \exp(\log p(\theta_i) + \log p(y|\theta_i)) / Z,$$

where the normalization constant is given by $Z = \int_{\Theta} \exp(\log p(\theta) + \log p(y|\theta)) d\theta$ and approximated using some numerical integration scheme based on the selected grid.

We note that the gridding approach can provide a reasonable approximation to the posterior, but it does not provide a straightforward method of sampling from this posterior.

3.3. Variational inference

Variational inference methods (Rezende and Mohamed, 2015; Beck, 1976; Guo et al., 2017; Hoffman et al., 2013; Blei et al., 2017) are a powerful and versatile class of techniques used in probabilistic and Bayesian modeling. They have an advantage of being very efficient, in particular for high-dimensional problems, compared to Markov chain Monte Carlo-based methods (Metropolis et al., 1953; Robert and Casella 2005; Robert and Casella, 2011). By formulating the posterior inference as an optimization problem, variational inference methods aim to approximate the Bayesian posterior with a tractable parametric distribution. Only the parameters of the parametric distribution are learned, rather than the typically intractable Bayesian posterior. An optimization problem is formed such that the parameters of the distribution are optimized by minimizing the Kullback–Leibler (KL) divergence (Cove and Thomas 2006) between the variational approximation and the true posterior. This optimization problem is formulated as minimizing the evidence lower bound (ELBO).

Variational inference assumes a parametric distribution $q_{\phi}(\theta|y)$ which is used to approximate the true posterior distribution $p(\theta|y)$. The goal of the variational inference is to minimize the discrepancy between the true and approximate distributions by optimizing the distribution parameters. To achieve this, a measure of discrepancy or distance between two distributions, called the KL divergence, is minimized between the two distributions. The KL divergence is given by

$$\text{KL}[q_{\phi}(\theta|y) \| p(\theta|y)] = \mathbb{E}_{\theta \sim q_{\phi}(\theta|y)} \left[\log \left(\frac{q_{\phi}(\theta|y)}{p(\theta|y)} \right) \right]. \quad (20)$$

Variational inference seeks to minimize this by solving the optimization problem

$$\phi^* = \arg \min_{\phi} \text{KL}[q_{\phi}(\theta|y) \| p(\theta|y)]. \quad (21)$$

Using Bayes' rule and defining the evidence lower bound (ELBO) $\mathcal{L}(q)$ as

$$\mathcal{L}(\theta, \phi) = \mathbb{E}_Q \left[\log \left(\frac{p(\theta, y)}{q_{\phi}(\theta|y)} \right) \right] = \mathbb{E}_Q[\log p(\theta, y)] - \mathbb{E}_Q[\log q_{\phi}(\theta|y)], \quad (22)$$

the KL divergence can be written as a function of the ELBO by

$$\text{KL}[q_{\phi}(\theta|y) \| p(\theta|y)] = -\mathcal{L}(\theta, \phi) + \log(p(y)). \quad (23)$$

The optimization problem of Eq. (21) is therefore equivalent to minimizing the negative ELBO:

$$\phi^* = \arg \min_{\phi} -\mathcal{L}(\theta, \phi)$$

3.3.1. Gaussian variational inference

Fixed-form variational inference uses simple parametric distributions as the variational posterior. In our experiments, we choose an independent Gaussian distribution $q_\phi(\theta|y) = \mathcal{N}(\mu_\phi, \Sigma_\phi)$ for the fixed form such that the parameters $\phi = [\mu_\phi, \Sigma_\phi]$ to be optimized consist of the mean $\mu_\phi \in \mathbb{R}^m$ and covariance $\Sigma_\phi \in \mathbb{R}^{m \times m}$ matrix. Note that the assumption of independence limits the covariance matrix to m trainable parameters and $2m$ total trainable parameters.

Gradient-free optimization techniques such as the Nelder–Mead algorithm (Nelder. and Mead 1965) scale poorly in high-dimensional problems and are often far less efficient than gradient-based optimization algorithms. This technique, therefore, requires the use of gradients to be effective, and gradients can be computed for the baseline model following the discussion in Section 3.4. We note that gradient-based algorithms are often trapped in local minima due to the often non-convex nature of general objective functions. However, we believe that the efficiency gains of gradient-based algorithms far outweigh the potential of finding suboptimal local minimum solutions.

3.4. Gradients through ODE solve

Our subsequent analysis will employ inference and ML methods that require gradient information. Therefore, ensuring that the ODE solve is differentiable and having an efficient method of computing gradients is critical. To obtain gradients of the ODE forward solve with respect to model parameters, we employ an adjoint-based method known as neural ODE (Böttcher and Asikis, 2022; Chen et al., 2018; Linot et al., 2023) using any black-box ODE solver. This method computes gradients through the ODE forward solve efficiently. It is also available as part of existing ML libraries such as Pytorch (Paszke et al., 2019), which allows its integration with other ML tools.

Propagating gradients through the deposition onset criteria requires extra care. Before film deposition begins, the resistance and film thickness do not increase. Therefore, the following dynamical system is solved:

$$\frac{dh}{dt} = 0, \quad \frac{dR_{\text{film}}}{dt} = 0, \quad (24)$$

with V , j , and Q computed analytically according to the experiment type. When both $j > j_{\text{min}}$ and $Q > Q_{\text{min}}$, deposition begins and the model dynamics instantaneously switch to

$$\frac{dh}{dt} = C_v j, \quad \frac{dR_{\text{film}}}{dt} = \rho(j) \frac{dh}{dt}. \quad (25)$$

As predicting the time of deposition onset is critical for predicting film growth, gradients of the output with respect to the onset condition parameters j_{min} and Q_{min} must be computed. However, the instantaneous change in model constitutes an in-place operation (Paszke et al., 2019) and must be treated separately. We use ideas from an extension to neural ODE, which adds the ability to compute gradients through instantaneous event handling (Chen et al., 2021). Adding an additional *switch state* ζ to our model, we computationally solve the following equations instead of Eqs. (24) and (25):

$$\frac{dh}{dt} = \zeta C_v j, \quad \frac{dR_{\text{film}}}{dt} = \zeta \rho(j) \frac{dh}{dt}. \quad (26)$$

The initial switch state $\zeta(t=0) = 0$ is switched to $\zeta(t=t_e) = 1$ at the event time, defined as the time t_e such that either of the deposition onset criteria are met. This implementation detail allows computing the gradient of the forward solve with respect to the event time, and in turn with respect to the onset criteria parameters j_{min} and Q_{min} , using the framework from Chen et al. (2021).

3.5. Parameter identifiability of baseline model

The process of parameter inference on the baseline model provides some insight into the shortcomings of the model and results in proposed updates to the baseline model to address some of these shortcomings.




Figure 4. Negative log-likelihoods computed from simulated data on a VR experiment using the baseline model with experimental conditions $V_R = 0.125$, $\sigma = 0.14$, $-\log C_v = 8.5$, $Q_{\min} = 100.0$, and $j_{\min} = 1.5$.

One of such shortcomings is parameter unidentifiability of the deposition onset criteria parameters, Q_{\min} and j_{\min} . The nature of this double criteria for the onset of deposition creates situations in which one parameter or the other may not be identifiable. For example, consider the case in which the time t_j at which $j > j_{\min}$ is smaller than the time t_Q at which $Q > Q_{\min}$, such that $t_j < t_Q$. Changing j_{\min} over some range will thus not have any effect on the baseline model output because *both* conditions $j > j_{\min}$ and $Q > Q_{\min}$ must be satisfied for deposition to begin. In this case, only Q_{\min} controls the deposition onset time and data may be uninformative about j_{\min} . However, the opposite can occur and there also exist cases in which data may be uninformative about Q_{\min} .

To illustrate this identifiability issue more clearly, we generate artificial data from the baseline forward model by taking the “true” parameters to be $[-\log C_v, Q_{\min}, j_{\min}] = [7.0, 150.0, 1.0]$. Other experimental parameters are given in Table 1. A total of 10 samples are obtained by simulating a VR experiment with a small amount of noise added. Figure 4 illustrates the negative log-likelihood of this data for each parameter to be inferred assuming that the other two parameters are fixed to their “true” values. For the first two parameters, a minimum exists at the “true” values, indicating that these parameters can be accurately inferred from the simulated data. However, the negative log-likelihood of the minimum current threshold parameter j_{\min} is flat over some region, indicating that the model is not effected by changes in j_{\min} over this region. If the “true” parameter value lies anywhere in the region in which the derivative of the log-likelihood is zero, that parameter is unidentifiable on that region. We use the term “unidentifiable” to mean that the first derivative of the likelihood is zero over some region in parameter space.

For VR experiments, there exists a set of boundaries in parameter space for which either j_{\min} or Q_{\min} will be unidentifiable, or both will be identifiable in baseline model. These boundaries change with the conditions of the VR experiment; gathering data from additional experiments could provide information on a parameter that is not informed by data from a different experiment.

Consider a case in which data exists such that $t_Q < t_j$. This means that t_j will control the deposition onset time. As $j(t) > 0 \forall t > 0$, then $Q(t > t_j) > Q(t_j) > Q_{\min}$. Thus changing Q_{\min} on a range $0 \leq Q_{\min} < Q(t_j)$ will have no effect on the output of the baseline model and the data will be uninformative about Q_{\min} on this range. Next consider a case in which data exists such that $t_j < t_Q$. Now t_Q will control the deposition onset time. However, $j(t > t_Q) > j(t_Q)$ is *not* guaranteed, which can be easily seen by taking the time derivative of Eq. (13). If the VR $V_R < \rho(j)C_v j$, then $dj/dt < 0$ and it is not guaranteed that $j(t > t_Q) > j(t_Q)$. If $j(t > t_Q) < j(t_Q)$, then it is possible for $j(t > t_Q) < j_{\min}$ and deposition stops, followed by an increase in current, restarting deposition, and the cycle repeats. Ultimately, this indicates that the parameter j_{\min} will have an influence on the baseline model output, and data will be informative about j_{\min} . Thus, if $t_j < t_Q$, data may or may not be informative about the parameter j_{\min} , depending on the experimental conditions.

Two identifiability regions corresponding to different experiments are computed empirically and illustrated in Figure 5. This regions are computed by simulating the baseline model according to the experimental conditions for a set of discretized points in the j_{\min}, Q_{\min} space. For each simulation,




Figure 5. Identifiability regions of the baseline model for two different experimental conditions. The log-likelihood on the simulated experimental data will be constant in the purple and cyan regions, indicating that little information is gained about j_{\min} if the true value lies in the purple region or Q_{\min} if the true value lies in the cyan region. Note: the “stepping” behavior observed in the identifiability boundaries here are a product of discretizing the j_{\min} and Q_{\min} domains, but the boundaries are in fact smooth.

identifiability is checked by checking t_Q and t_j . Purple regions indicate that data from the experiment are not informative about j_{\min} if the true value of j_{\min} and Q_{\min} lie in the region. In other words, $\partial\mathcal{L}/\partial j_{\min}=0$ on the purple region. Cyan regions indicate that data from the experiment is not informative about Q_{\min} if the true values of j_{\min} and Q_{\min} lie in the region, or $\partial\mathcal{L}/\partial j_{\min}=0$. Yellow regions indicate that the experiment can inform both j_{\min} and Q_{\min} . The boundary of the cyan region can be computed analytically, but the other is nontrivial and computed empirically. The times at which the deposition onset criteria are met in the VR experiment for the baseline model are given by

$$t_j = \frac{2Q_{\min}}{\sigma V_R}(\sigma R_0 + L), \quad t_Q = \left[\frac{2j_{\min}}{\sigma V_R}(\sigma R_0 + L) \right]^{1/2}. \tag{27}$$

Setting Eq. (27) equal, we obtain a closed form solution to the Q_{\min} identifiability boundary as

$$j_{\min} = \left[\frac{2Q_{\min}\sigma V_R}{\sigma R_0 + L} \right]^{1/2}, \tag{28}$$

which has been validated against the empirically computed boundaries shown in Figure 5. We note that the log-likelihoods in Figures. 4b and 7c correspond to the identifiability plane in Figure 5a. The log-likelihoods of Q_{\min} and j_{\min} have first derivative zero over some region as predicted by the identifiability boundaries. The true values of j_{\min} and Q_{\min} are such that j_{\min} is not informed by the simulated experimental data.

If the experimental data does not inform one of the parameters, then that parameter does not influence the output of our baseline model, and it cannot be accurately inferred. However, as shown by Figure 5, the identifiability boundaries change with the type of experiment. This behavior could cause very poor prediction results. Suppose none of the experimental data is informative about j_{\min} , and inference is performed on the model parameters. Using the model in prediction could result in poor performance, especially if predictions are made for an experimental condition in which j_{\min} does influence the output of the model.

It is also necessary to infer robust posteriors in this case. Suppose Gaussian variational inference is selected as the inference method. A Gaussian variational posterior will be computed for each of the parameters to be inferred, providing a unimodal distribution for each. This can be misleading about the

Bayesian posterior distribution of the parameter and result in poor uncertainty quantification during prediction.

We explore two approaches toward the aim of improving model predictions by improving the form of the baseline model. First, we update the model based on insight from the shortcomings observed during inference, as described in Section 4.1. We then pursue a different approach in which model augmentations are learning using machine-learning tools to augment the baseline model with previously unmodeled dynamics, as discussed in Section 4.2.

3.6. Inference on baseline model


In this section, we demonstrate inference results using Gaussian variational inference on the baseline model. The purpose of this experiment is to illustrate the shortcomings of the model form itself, not the selected method of inference. Inference is performed using simulated data from the baseline model, and we demonstrate that the posterior predictive results in good prediction performance for some experimental conditions, but poor performance on others.

Data are generated by simulating the baseline model 10 times up to a final time $T=250s$ for a VR experiment with $V_R=1.0$, $\log C_V=-8$, $j_{\min}=1.5$, $\sigma=0.14$, and $Q_{\min}=100$, and add Gaussian random noise $\epsilon \sim \mathcal{N}(0, \eta^2)$, where $\eta=0.01$, at each time step to simulate measurement noise. Only data from current measurements are considered in the inference process. We then perform Gaussian VI for parameters $\log C_V$, j_{\min} , and Q_{\min} on the baseline model.

After learning the variational posterior q_ϕ , samples are drawn and used to simulate the baseline model to obtain samples from the posterior predictive. Ten of these samples along with the mean are shown in Figure 6a. For the experimental configuration on which the data are generated, accurate and low variance prediction is observed. However, we then use the variational posterior distribution to predict on a VR experiment in which all parameters are the same except the voltage ramp $V_R=0.125$. Figure 6b shows that prediction performance is very poor for this experiment.

The reason for this stems from the identifiability issues previously discussed in Section 3.5. In our experiment, we assume that the ‘true’ value of j_{\min} is 1.5. However, in the experiment which we gather data from, j_{\min} is unidentifiable in the baseline model. Thus, the inference results in a variational posterior distribution for j_{\min} of $q(j_{\min}|\mathcal{D})=\mathcal{N}(2.66, 0.01)$, which is a low variance but poor estimate of the true parameter value. This posterior distribution is then used for prediction in an experiment in which j_{\min} *does* have an effect on model output, and because the posterior is not accurate, prediction is also inaccurate.

On top of identifiability issues potentially resulting in poor prediction performance, the parameter Q_{\min} is inconsistent across different experiment types. To illustrate this issue, Gaussian VI is performed on the



(a) VR experiment data, $V_R = 1.0$, $\sigma = 0.14$, $-\log C_V = 8$, $j_{\min} = 1.5$, $Q_{\min} = 100$

(b) VR experiment prediction, $V_R = 0.125$, $\sigma = 0.14$, $-\log C_V = 8$, $j_{\min} = 1.5$, $Q_{\min} = 100$

Figure 6. Posterior predictive results after performing Gaussian VI on data from a simulated VR experiment. The posterior predictive results in accurate simulations on the data (a), but poor predictions for other experiments (b). This is caused by unidentifiable j_{\min} in the data.

parameters C_v, j_{\min} , and Q_{\min} twice—first using real experimental data from only the VR experiments and again using only data from the CC experiments. The MAP of the variational posterior for each distribution are very different—in particular for Q_{\min} . Using only VR experimental data, the MAP of the variational posterior is $Q_{\min} \approx 261$; however, using only CC experimental data, the MAP is located at $Q_{\min} \approx 101$. This is indicative of Q_{\min} being problematic in allowing the baseline model to accurately predict experimental data with different boundary conditions. Our natural conclusion is that the model is incorrect, and a root cause may lie in a constant Q_{\min} . Rather, Q_{\min} may be a function of the type of experiment being performed rather than a constant to be inferred. In Section 4.1, physically intuitive model updates are introduced with the aim of alleviating the identifiability concerns of the baseline model and improving the modeling of the minimum charge criterion.

4. Model updates

The baseline model was shown to exhibit identifiability as well as generalization deficiencies in Section 3. In this section, we aim to improve the prediction accuracy and generalizability of the baseline model. First, modifications are made based on the inference results of the baseline model to aid in improving parameter identifiability and minimum charge criterion modeling.

An alternative approach to improving the baseline model is also investigated in which machine-learning augmentations are introduced to model system dynamics which are absent from the baseline model. These augmentations are introduced with an emphasis on interpretability of the augmented model while allowing for greater flexibility in model expressiveness.

4.1. Inference-informed modifications

Based on the inference experiments of Section 3, model updates are proposed to alleviate the observed inadequacies during prediction. The issue of identifiability arises in the baseline model due to the double conditional statement that film deposition begins only if $j > j_{\min}$ and $Q > Q_{\min}$. This conditional statement, in particular $j > j_{\min}$, creates non-physical, discontinuous behavior of the model. The film growth rate given by Eq. (24) is exactly zero until the conditional statement is true. Assuming that $j_{\min} > 0$, the film growth rate will instantaneously increase, and a sharp discontinuity in the film thickness growth occurs. The model also does not accurately model the cases in which both onset criteria are met, but the minimum current condition is no longer met at a later time. This behavior is one of the reasons that prediction is observed to be inaccurate in Figure 6b. We therefore propose a model update to create a model in which the dynamics are continuous which also allows for film dissolution by replacing the film thickness dynamics in Eq. (25) with

$$\frac{dh}{dt} = C_v(j_n - j_{\min}) \text{ for } Q > Q_{\min}, \text{ s.t. } h \geq 0. \quad (29)$$

This also gives the benefit of j_{\min} being identifiable for all experimental configurations.

With j_{\min} now identifiable, the expressiveness of the parameter is investigated to improve generalizability. Assuming that the minimum charge criterion is related to the concentration of OH^- present in the bath, C illustrates that Q_{\min} is not constant across experiment types, but depends on a different constant K . It is shown in C that Q_{\min} is a function of this constant, and the function differs between the VR and CC experiments. For VR (Eq. (30)) and CC (Eq. (31)) experiments, this function is given by

$$Q_{\min} = \left(\frac{81}{128\beta} \right)^{1/3} K^{4/3} \quad (30)$$

$$Q_{\min} = \frac{K^2}{j_0}, \quad (31)$$

where $\beta = \sigma V_R / (\sigma R_0 + L)$.




Figure 7. Negative log-likelihoods computed from simulated data on a VR experiment using the inference-informed model with experimental conditions $V_R = 0.125$, $\sigma = 0.14$, $-\log C_v = 8.5$, $Q_{\min} = 100.0$ ($K = 23.2$), and $j_{\min} = 1.5$.

The updated film thickness dynamics of Eq. (29) along with introducing the parameter K used to compute the minimum charge criterion constitute an updated model which we dub the “inference-informed” model.

Performing the same exercise to visualize the negative log-likelihood as in Figure 4, we visualize the negative log-likelihood of the inference-informed model on artificial data to illustrate that all parameters are now identifiable. In this case, we simulate the data from the same experimental configuration as Figure 4, which is a VR experiment with $V_R = 0.125$, $\sigma = 0.14$, $-\log C_v = 8.5$, $Q_{\min} = 100.0$, and $j_{\min} = 1.0$. However, the parameter K is used instead of Q_{\min} ; thus, the value of K is found using the relationship in Eq. (30), obtaining $K = 23.2$. The negative log-likelihoods of $-\log C_v$, K , and j_{\min} are illustrated in Figure 7, and all parameters exhibit a global minimum at the true parameter values, indicating that all are now identifiable.

4.1.1. Model comparisons

The baseline model and the inference-informed model are directly compared by performing Bayesian inference using the gridding approach discussed in Section 3.2. We then predict for each model based on the maximum a posteriori and compute the negative log-likelihood of the data. Lower values of the negative log-likelihood at the MAP correspond to a model which better fits the experimental data.

The MAP of the approximated posterior using the gridding approach assuming the baseline model is located at $\theta_{\text{MAP}} = [-\log C_v, Q_{\min}, j_{\min}] = [7.58, 173.5, 0.0]$ with a negative log-likelihood at the MAP of 3.35×10^7 . Assuming the updated model, the MAP is located at $\theta_{\text{MAP}} = [-\log C_v, K, j_{\min}] = [7.32, 44.2, 0.63]$ while the negative log-likelihood is 2.65×10^7 . This suggests that the inference-informed model performs better than the baseline model by having the potential to represent the experimental data more accurately.

Each model prediction at the map is compared to the experimental data in Figure 8. The current predictions are shown for two experiments: VR with $V_R = 1.0 \text{ V/s}$ and CC with $j_0 = 7.5 \text{ mA}$. Predictions for all experiments on current, resistance, and thickness data are included in D. Prediction of the inference-informed model for CC experiments exhibits significant improvement over the baseline model, likely due to the improved parameterization of the minimum charge criterion Q_{\min} . However, there is clearly some physical behavior which exists in the VR experiments which is not present in our model. For instance, there are two current “peaks” in experimental data for VR experiments, but only one is possible in model predictions. In addition, although the thickness prediction of the inference-informed model is closer to the experimental data than the baseline model, Figure 9 illustrates that neither are particularly accurate. Ultimately, thickness prediction is the most important quantity of interest. We thus turn toward improving the model through augmenting the model forming and learning the augmentations via a machine-learning based framework.

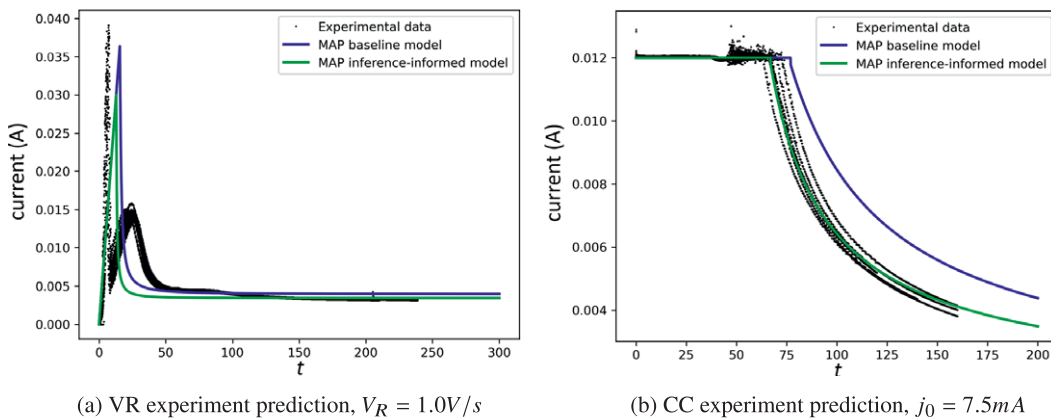


Figure 8. Comparisons between current prediction on the baseline model and inference-informed model at the MAP for each on (a) VR experiment with $V_R = 1.0\text{V/s}$ and (b) CC experiment with $j_0 = 7.5\text{mA}$.

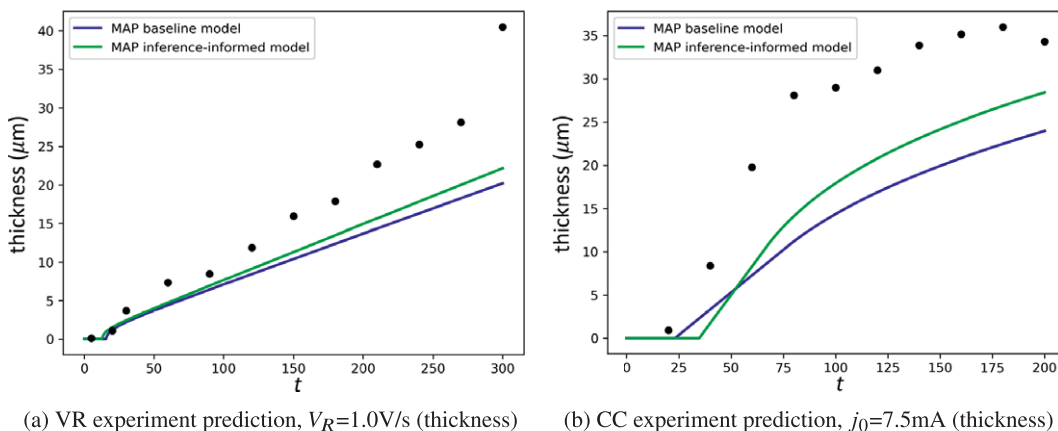


Figure 9. Comparisons between film thickness prediction on the baseline model and inference-informed model at the MAP for each.

4.2. Machine-learning augmentations

Modifying the model using insights from the inference process aids in understanding some shortcomings due to model form but results in limited improvement to predictive performance. Additional physical behavior is missing from the model itself, which results in poor performance even when the optimal model parameters are identified. There are two main features absent from the baseline model: the presence of two peaks in the current dynamics for the VR experiments and smooth behavior of those peaks. Unstable behavior in data for CC experiments observed just before the drop in current is due to the current controller in the experiments, which we do not account for in our model.

To incorporate the two missing physical features in our model, we turn to machine learning with an emphasis on interpretability. In particular, we augment the right hand side of the dynamics model with parameterized neural networks and train the augmentations using NeuralODE (Chen et al., 2018). This can be used effectively to utilize adjoint-based gradient computation alongside standard machine-learning pipelines to learn flexible augmentations in the right hand side of a dynamical system.

We first attempt to characterize the underlying physics of each of the two peaks present in the experimental current data for VR experiments. The second peak corresponds to the onset of deposition, which is previously modeled in the baseline and inference-informed models. As film deposition begins,

resistance increases and current decreases. However, the baseline model assumes an instantaneous onset of film deposition which results in a discontinuous and physically impossible film growth rate. Experimental data illustrate smooth transitions to film growth in current data, further supporting the need for additional modeling. This smooth transition is likely the result of portions of the surface being coated in a non-uniform manner, resulting in only a fraction of the material surface being coated for a small time. We propose an augmentation to the baseline model, which has the additional benefit of modeling the deposition onset time without the need of threshold parameters. We propose multiplying the right hand side of Eq. (25) by a learnable term $g_\phi(V, Q)$ which varies smoothly between zero and one. This term is dubbed the “coverage fraction model,” and it is a function of voltage and charge only, representing the coverage fraction of the film on the material. It is found empirically that using both voltage and charge as inputs to the model results in the best predictive performance. It is a function of charge due to the dependence of the deposition onset time of the baseline model on charge, and it is a function of voltage due to a change in the time derivative of voltage across experiment types.

The first peak present in the current data for VR experiments is caused by a phenomenon ignored by the baseline model altogether. This peak may be caused by an oxygen-reduction reaction (Nagai et al., 2012), which occurs at the anode as the voltage increases. The relationship between electric potential and current in redox reactions is described by the (Dickinson and Wain, 2020) Eq. (11) of the form

$$j = j_0 [\exp(aV) - \exp(bV)], \tag{32}$$

where the parameters a and b depend on many physically relevant parameters. However, the particular form is not relevant to the discussion here as we aim to learn this component of the model while keeping the general form to ensure that the model is physically interpretable. We assume that the combined effect of this redox reaction results in a resistance “source” such that $j = c_1 \exp(c_2 V)$, based on the form of Eq. (32).

After some critical point is reached (which we do not explicitly model, but likely corresponds to some minimum charge criterion), we assume that the OH^- starts being diffused at some rate (Nagai et al., 2012), which decreases the OH^- concentration. As the redox reaction continues, we assume that the combined effect of these two reactions results in a different exponent such that $j = c_3 \exp(c_4 V)$, where $c_4 < 0$.

To augment the VR experiment model with this behavior of switching between two different exponents resulting from a combination of reactions, we use an exponential function in which the argument includes a hyperbolic tangent function which can vary smoothly between two different values. We thus assume an augmented model of the form

$$j = \frac{\sigma V}{\sigma R_{\text{film}} + L} + c_1 \exp(c_2 V f_\theta(V, Q)), \quad f_\theta(V, Q) = \tanh\left(-\hat{f}_\theta(V, Q)\right), \tag{33}$$

The augmented model is finally expressed by

$$f_\theta(V, Q) = \tanh\left(-\hat{f}_\theta(V, Q)\right), \quad g_\phi(V, Q) = \frac{1}{1 + c_3 \exp(-c_4 \hat{g}_\phi(V, Q))} \tag{34}$$

$$\frac{dR_{\text{film}}}{dt} = g_\phi(V) \rho(j) j C_v, \quad j = c_1 (\exp(c_2 V f_\theta(V))) + \frac{\sigma V}{\sigma R_{\text{film}} + L}, \tag{35}$$

for VR experiments and

$$g_\phi(V, Q) = \frac{1}{1 + c_3 \exp(-c_4 \hat{g}_\phi(V, Q))}, \quad \frac{dR_{\text{film}}}{dt} = g_\phi(V) \rho(j) j C_v, \tag{36}$$

for CC experiments. Note that the relationship between voltage, current, and resistance are unchanged from the baseline model in the CC experiments even with the machine-learning augmentations introduced to the model. The model augmentation functions \hat{g}_ϕ and \hat{f}_θ are parameterized by FNNs with four layers, eight nodes per layer, and ReLU activation functions. These are learned along with the constants c_1, c_2, c_3

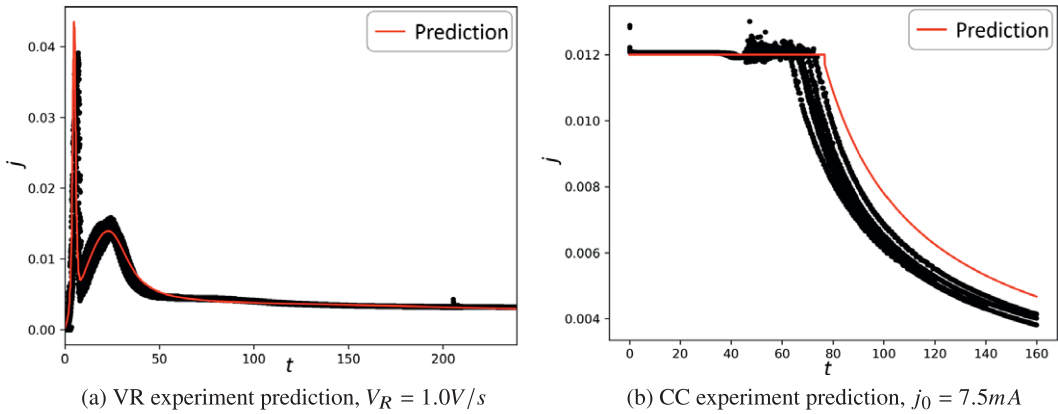


Figure 10. Current prediction on the machine-learning augmented model trained with the first peak model.

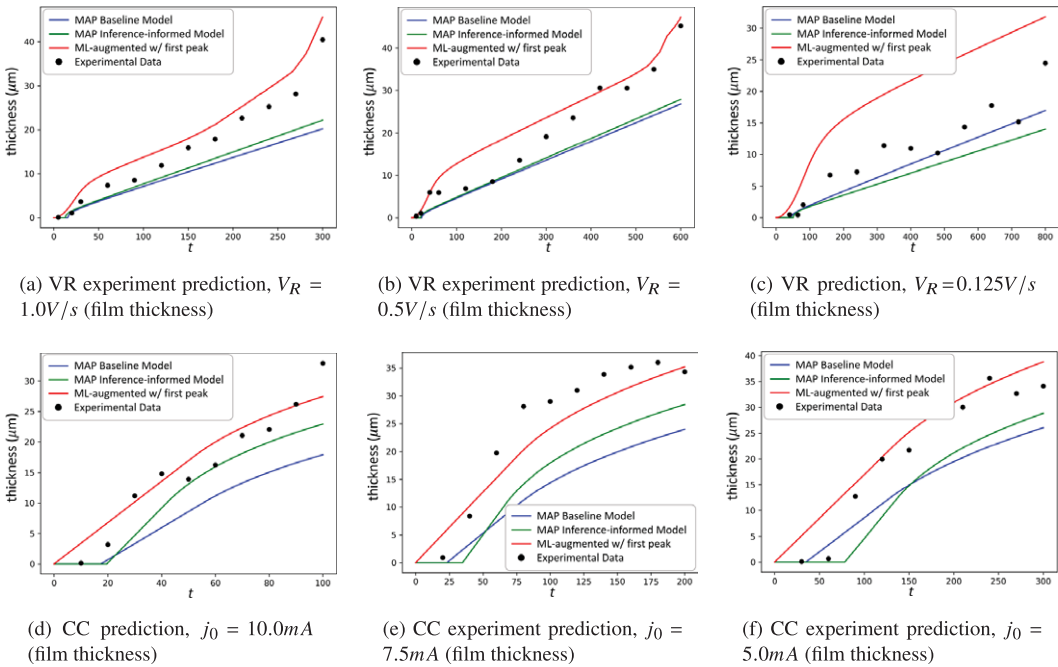


Figure 11. Comparisons between film thickness prediction on the baseline model, inference-informed model, and ML-augmented model with first peak.

and c_4 using the NeuralODE framework using only current data from the three VR experiments. Thus, applying the model on the CC experiments is purely prediction, as none of the data is seen during training. We learn the parameters $\theta, \phi, C_v, \sigma, c_1, c_2, c_3,$ and c_4 in our experiments and show prediction results. Figure 10 illustrates the results of the final trained model, better capturing the “double-peaked” nature of the current in experimental data for VR experiments. Here, we show current predictions for only two experimental configurations, but current and resistance predictions for all other experimental configurations are provided in E.

Furthermore, thickness predictions for all experimental configurations are illustrated in Figure 11 along with predictions from the baseline and inference-informed models for comparison. Film thickness




Figure 12. Current prediction on the machine-learning augmented model trained without the first peak model, shown for (a) VR experiment with $V_R = 1.0\text{V/s}$ and (b) CC experiment with $j_0 = 7.5\text{mA}$.

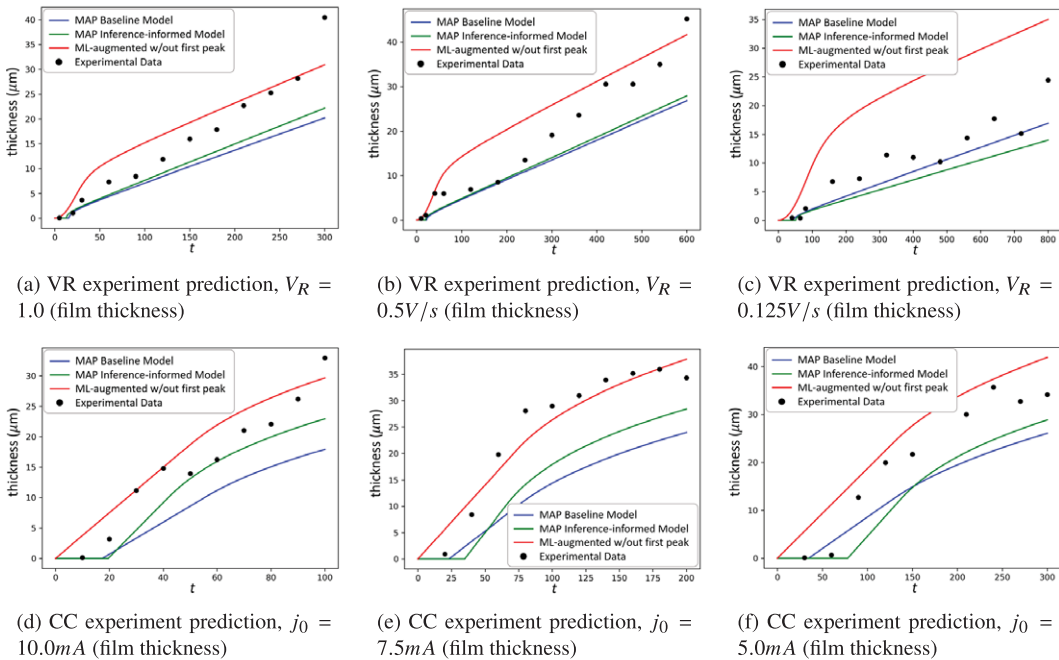


Figure 13. Comparisons between film thickness prediction on the baseline model, inference-informed model, and ML-augmented model without first peak.

predictions show an improvement over the baseline model in all cases except for a single experimental configuration: VR experiments with $V_R = 0.125$. We expect that this is due to some unmodeled behavior in the low-voltage regime which is difficult to capture. However, in CC experiments and larger voltage values in VR experiments, our augmentations provide significant improvement to model predictions.

Learning the first peak model function f_θ in the neural ODE framework is significantly more expensive computationally than training only the coverage fraction model g_ϕ due to the large gradient change requiring a finely discretized time step to adequately resolve. In addition, the first peak model performs notably poorly for the VR experiment in the low-voltage ramp regime. We thus remove the “first peak” model and retrain only the augmentation function $g_\phi(V, Q)$ by masking out the experimental data

corresponding to the first peak in current for VR experiments. This results in more efficient training and thickness prediction while providing quite similar predictions when the first peak model is included. These results show that modeling the dynamics of the first peak may unnecessarily decrease computational efficiency while still providing more accurate thickness prediction, which is ultimately the goal. Current predictions for two experiments (one VR and one CC) are shown in Figure 12, and thickness predictions are shown and compared to all other models presented in Figure 13. These results again illustrate a significant improvement in thickness prediction without the need for including the more expensive first peak model, indicating that modeling such behavior may be largely unnecessary for predicting thickness. However, prediction in the low voltage ramp regime is actually worse than with the first peak model included, further supporting our hypothesis that there exist additional unmodeled dynamics in such cases. Again, all other predictions for current and resistance using our model augmentations without the first peak model included are provided in F.

5. Conclusions

We presented a comprehensive framework to introduce adaptable yet interpretable machine-learning-based model augmentations. Our initial objective of parameter inference unveiled significant limitations in the existing baseline model, necessitating the integration of diverse model refinements aimed at rectifying these deficiencies.

Even with such enhancements, however, certain physical behaviors inherent in experimental data were not fully replicated by the models. Notably, instances such as the emergence of two peaks in current data during VR experiments and the smooth evolution of current data remained elusive. To address these disparities, we devised physically meaningful augmentations, seamlessly integrated them into the model, and harnessed a neural ODE framework for implementing learnability. This approach to augmentation construction maintains interpretability while harnessing the expressive power of neural networks to allow enhanced adaptability.

The integration of these augmentations yields models that demonstrate improved predictive accuracy when compared against empirical data. Nonetheless, limitations persist. One key challenge lies in the computational costs associated with modeling the “first peak.” While we found that accurate film thickness predictions could still be achieved without modeling this first peak, its omission leads to a reduction in computational burdens. Despite the reduced offline computational cost, especially when the first peak model is omitted, there remain areas for improvement in balancing accuracy and efficiency, particularly for real-time applications or online model improvement.

Future work could explore alternative augmentation methods or hybrid modeling techniques to better capture unmodeled phenomena without increasing computational demands. Although the ML-enhanced models required some offline computational cost, the online prediction cost remains nearly identical to the baseline model due to the simple and efficient FNN-based model augmentations. Offline optimization converged within a few hours on a single GPU with the first peak model and within an hour without it. Future efforts might focus on reducing this offline burden further while preserving the same degree of accuracy.

Integral to our methodology is an in-depth procedure of evaluating the model’s behavior prior to augmentation. This diagnostic analysis is pivotal, as it grants insights into the limitations of the model. By first grasping these limitations, we gather information on how to design augmentations that retain physical interpretability, while simultaneously fostering flexibility and removing the necessity for manual crafting and fine tuning of model structures. This practice holds immense significance for machine-learning-driven dynamical system understanding and for broader applications in physics-based domains. It facilitates augmentation strategies that preserve interpretability, thereby opening avenues for model improvements.

Future research should focus on extending this methodology to more complex or transient/chaotic dynamic systems where interpretability could become a trade-off with predictive power. In addition,

future studies might investigate how well the proposed framework generalizes to additional unseen conditions and environments, ensuring its robustness across varied applications.

We acknowledge that other problem settings may introduce different intricacies; nevertheless, the systematic process outlined in this study can guide future applications in different problems. This study offers a variety of insights, equipping researchers with tools to assess and elevate the forms of model representations for diverse dynamical systems. Ultimately, our work fosters a bridge between machine learning and physics, improving the modeling process while revealing the underlying system dynamics, and setting the stage for future exploration into even more efficient and adaptable approaches.

Data availability statement. All data in this study are available only through approval by Ford Motor Company. The code for our experiments is available at <https://github.com/christian-jacobsen/NormalizingFlowsInference>.

Author contribution. Conceptualization: K.D.; W.L.; M.A.; C.J. Methodology: C.J; M.K. Data curation: W.L.;M.K. Data visualisation: C.J. Writing original draft: C.J. All authors approved the final submitted draft.

Competing interest. None

Funding statement. This work was supported at the University of Michigan by Ford Motor Company under the grant “Hybrid Physics-Machine Learning Models for Electro-deposition.”

Ethical standard. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

References

- Ahmed SE, Pawar S, San O, Rasheed A, Iliescu T and Noack BR (2021) On closures for reduced order models—A spectrum of first-principle to machine-learned avenues. *Physics of Fluids* 33 (9), 091301.
- Beck F (1976) Fundamental aspects of electrodeposition of paint. *Progress in Organic Coatings* 4 (1), 1–60.
- Blei DM, Kucukelbir A and McAuliffe JD (2017) Variational inference: A review for statisticians. *Journal of the American Statistical Association* 112 (518), 859–877
- Böttcher L and Asikis T (2022) Near-optimal control of dynamical systems with neural ordinary differential equations. *Machine Learning: Science and Technology* 3 (4), 045004.
- Boyd DW and Zwack RR (1996) Predicting electrocoat efficiency: A simple data-based model for predicting electrocoat usage. *Progress in Organic Coatings* 27 (1), 25–32 Proceedings of the 20th International Conference in Organic Coatings Science and Technology.
- Brunton SL, Proctor JL and Kutz JN (2016) Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences* 113 (15), 3932–3937.
- Chen, R. T. Q., Amos, B., and Nickel, M. (2021). Learning neural event functions for ordinary differential equations. In *International Conference on Learning Representations*.
- Chen RTQ, Rubanova Y, Bettencourt J and Duvenaud DK (2018) Neural ordinary differential equations. In Bengio S, Wallach H, Larochelle H, Grauman K, Cesa-Bianchi N and Garnett R (eds), *Advances in Neural Information Processing Systems*, Vol. 31. Curran Associates, Inc.
- Cheng S, Quilodrán-Casas C, Ouala S, Farchi A, Liu C, Tandeo P, Fablet R, Lucor D, Iooss B, Brajard J, Xiao D, Janjic T, Ding W, Guo Y, Carrassi A, Bocquet M and Arcucci R (2023) Machine learning with data assimilation and uncertainty quantification for dynamical systems: A review. *IEEE/CAA Journal of Automatica Sinica* 10 (6), 1361–1387.
- Cover TM and Thomas JA (2006) *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. USA: Wiley-Interscience. 18–30.
- Dickinson EJ and Wain AJ (2020) The butler-volmer equation in electrochemical theory: Origins, value, and practical application. *Journal of Electroanalytical Chemistry* 872, 114145 Dr. Richard Compton 65th birthday Special issue.
- Duraisamy, K. (2021). Variational encoders and autoencoders : Information-theoretic inference and closed-form solutions.
- Ellwood K, Tardiff JL, Gray LJ, Gaffney P, Braslaw J, Moldekleiv K and Halvorsen A (2009) Development of a full vehicle electrocoat paint simulation tool. *SAE International Journal of Materials and Manufacturing* 2 (1), 234–240.
- Guo, F., Wang, X., Fan, K., Broderick, T., and Dunson, D. B. (2017). Boosting variational inference.
- Hasegawa K, Fukami K, Murata T and Fukagata K (2020) Machine-learning-based reduced-order modeling for unsteady flows around bluff bodies of various shapes. *Theoretical and Computational Fluid Dynamics* 34 (4), 367–383.
- Hoffman MD, Blei DM, Wang C and Paisley J (2013) Stochastic variational inference. *Journal of Machine Learning Research* 14 (40), 1303–1347.
- Linot AJ, Burby JW, Tang Q, Balaprakash P, Graham MD and Maulik R (2023) Stabilized neural ordinary differential equations for long-time forecasting of dynamical systems. *Journal of Computational Physics* 474, 111838.
- Liu W and Kevin Ellwood FMC (2017) *Simulation of the Full Vehicle Electrodeposition for Paint Virtual Manufacturing*. SAE World Congress.

- Marlar T, Trainor A, Liu W, Ellwood K, Okerberg B, Padash F and Harb JN** (2020) Effect of convection on initial deposition during electrocoating of galvanized steel. *Journal of the Electrochemical Society* 167 (10), 103502.
- Metropolis N, Rosenbluth AW, Rosenbluth MN, Teller AH and Teller E** (1953) Equation of state calculations by fast computing machines. *The Journal of Chemical Physics* 21 (6), 1087–1092.
- Mišković-Stanković VB** (2002) The mechanism of cathodic electrodeposition of epoxy coatings and the corrosion behaviour of the electrodeposited coatings. *Journal of the Serbian Chemical Society* 67 (5), 305–324.
- Nagai H, Onishi Y and Amaya K** (2012) Exact paint resistance and deposition models for accurate numerical electrodeposition simulation. *Transactions of the Japan Society of Mechanical Engineers Series A* 78 (794), 1446–1461.
- Nelder JA and Mead R** (1965) A simplex method for function minimization. *The Computer Journal* 7 (4), 308–313.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S.** (2019). Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, 32 pages 8024–8035. Curran Associates, Inc.
- Pawar S, Rahman SM, Vaddireddy H, San O, Rasheed A and Vedula P** (2019) A deep learning enabler for nonintrusive reduced order modeling of fluid flows. *Physics of Fluids* 31 (8), 085101.
- Pierce PE** (1981) The physical chemistry of the cathodic electrodeposition process. *Journal of Coatings Technology* 53 (672), 52–67.
- Rastegar S, Ranjbar Z and Fakhropour G** (2008) A generalized model for the film growth during the constant voltage electrodeposition of resin dispersions. *Colloids and Surfaces A: Physicochemical and Engineering Aspects* 317 (1), 17–22.
- Rezende D and Mohamed S** (2015) Variational inference with normalizing flows. In Bach F and Blei D (eds), *Proceedings of the 32nd International Conference on Machine Learning, Volume 37 of Proceedings of Machine Learning Research*. Lille, France: PMLR, pp. 1530–1538.
- Robert CP and Casella G** (2005) *Monte Carlo Statistical Methods (Springer Texts in Statistics)*. Berlin, Heidelberg: Springer-Verlag. 79–318.
- Robert C and Casella G** (2011) A short history of markov chain Monte Carlo: Subjective recollections from incomplete data. *Statistical Science* 26 (1).
- Rudy SH, Brunton SL, Proctor JL and Kutz JN** (2017) Data-driven discovery of partial differential equations. *Science Advances* 3 (4), e1602614.
- San O and Maulik R** (2018) Neural network closures for nonlinear model order reduction. *Advances in Computational Mathematics* 44 (6), 1717–1750.
- Sand HJS** (1899) III. On the concentration at the electrodes in a solution, with special reference to the liberation of hydrogen by electrolysis of a mixture of copper sulphate and sulphuric acid. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 1 (1), 45–79.
- Tran, M.-N., Nguyen, T.-N., and Dao, V.-H.** (2021). A practical tutorial on variational bayes.

A. Voltage ramp

For the electric field, we solve a Poisson equation with Robin boundary condition on the interface bath/film and Dirichlet condition at the anode.

$$\sigma_{\text{bath}} \frac{\partial^2 \phi}{\partial x^2} = 0 \quad \text{in the bath} \quad (37)$$

$$\phi - R_{\text{film}} \sigma_{\text{bath}} \frac{\partial \phi}{\partial x} = 0 \quad \text{at the interface film-bath} \quad (38)$$

$$\phi_{\text{anode}}(t) = \phi_{t=0} + \phi_{\text{ramp}}(t) \quad \text{at the anode} \quad (39)$$

where σ_{bath} is the bath conductivity, j is the normal component of the current density, ϕ is the electrical potential, R_{film} is the film resistance, and h is the film thickness.

B. Constant current

For the electric field, we solve a Poisson equation with Robin boundary condition on the interface bath/film and Neumann condition at the anode.

$$\sigma_{\text{bath}} \frac{\partial^2 \phi}{\partial x^2} = 0 \quad \text{in the bath} \quad (40)$$

$$\phi - R_{\text{film}} \sigma_{\text{bath}} \frac{\partial \phi}{\partial x} = 0 \quad \text{at the interface film-bath} \quad (41)$$

$$\sigma_{\text{bath}} \frac{\partial \phi}{\partial x} = j_0 \quad \text{at the anode} \quad (42)$$

C. Evolution of the concentration of OH⁻

C.1. Solution of the diffusion equation

We consider the diffusion equation in the domain [0;L]

$$\frac{\partial u}{\partial t} = D \frac{\partial^2 u}{\partial x^2} \tag{43}$$

with the following boundary conditions:

$$\left. \frac{\partial u}{\partial x} \right|_{x=0} = -\frac{j(t)}{DF} = g(t) \text{ and } u(x=L, t) = h(t) \tag{44}$$

where u is the OH⁻ concentration, F is the Faraday constant, D is the diffusion coefficient, h , and g are time-dependent functions. The initial condition is defined by $u(x, 0) = u_0$, u_0 being the initial concentration.

Applying the Laplace transform

$$\mathcal{L} \left\{ \frac{\partial u}{\partial t} - D \frac{\partial^2 u}{\partial x^2} \right\} = \mathcal{L} \left\{ \frac{\partial u}{\partial t} \right\} - D \left\{ \frac{\partial^2 u}{\partial x^2} \right\} = s\bar{u} - u_0 - D\bar{u}_{xx} = 0 \tag{45}$$

The solution of the homogeneous equation $s\bar{u} - D\bar{u}_{xx} = 0$ is given by

$$\bar{u}_h(x, s) = C_1 \exp(r_1 x) + C_2 \exp(r_2 x) \tag{46}$$

where $r_1 = \sqrt{s/D}$ and $r_2 = -\sqrt{s/D}$

The specific solution, assuming a constant $\bar{u}_{sp.} = A$, is given by $\bar{u}_{sp.} = u_0/s$. Therefore, the solution has the following form

$$\bar{u}(x, s) = C_1 \exp(r_1 x) + C_2 \exp(r_2 x) + \frac{u_0}{s} \tag{47}$$

where C_1 and C_2 are functions derived using the boundary and initial conditions.

Using the Neumann boundary condition at $x = 0$ and the definition $\mathcal{L}\{g(t)\} = G(s)$, then

$$\bar{u}_x(x=0, s) = G(s) = C_1 r_1 + C_2 r_2. \tag{48}$$

Therefore, $C_1 = \frac{G(s) - C_2 r_2}{r_1} = \frac{G(s)}{r_1} + C_2$.

Hence,

$$\bar{u}(x, s) = \left(\frac{G(s)}{r_1} + C_2 \right) \exp(r_1 x) + C_2 \exp(r_2 x) + \frac{u_0}{s} \tag{49}$$

Using the Dirichlet boundary condition at $x = L$:

$$\bar{u}(x=L, s) = H(s) = \left(\frac{G(s)}{r_1} + C_2 \right) \exp(r_1 L) + C_2 \exp(r_2 L) + \frac{u_0}{s}, \tag{50}$$

the solution is given by

$$\bar{u}(x, s) = \left[\frac{G(s)}{r_1} + \frac{H(s) - \frac{u_0}{s} - \frac{G(s)}{r_1} \exp(r_1 L)}{\exp(r_1 L) + \exp(r_2 L)} \right] \exp(r_1 x) + \left[\frac{H(s) - \frac{u_0}{s} - \frac{G(s)}{r_1} \exp(r_1 L)}{\exp(r_1 L) + \exp(r_2 L)} \right] \exp(r_2 x) + \frac{u_0}{s} \tag{51}$$

If $L \rightarrow \infty$ and $h(t) = u_0$ (i.e. $H(s) = u_0/s$), the solution simplifies to

$$\bar{u}(x, s) = \left[\frac{G(s)}{r_1} - \frac{G(s)}{r_1} \right] \exp(r_1 x) + \left[-\frac{G(s)}{r_1} \right] \exp(r_2 x) + \frac{u_0}{s} \tag{52}$$

$$= -\frac{G(s)}{r_1} \exp(r_2 x) + \frac{u_0}{s} \tag{53}$$

Using the inverse Laplace transform yields

$$\mathcal{L}^{-1} \{ \bar{u}(x, s) \} = \mathcal{L}^{-1} \left\{ -G(s) \sqrt{\frac{D}{s}} \exp\left(-\sqrt{\frac{s}{D}} x\right) + \frac{u_0}{s} \right\} \tag{54}$$

$$= -\sqrt{D} \mathcal{L}^{-1} \left\{ G(s) \frac{1}{\sqrt{s}} \exp\left(-\frac{x}{\sqrt{D}} \sqrt{s}\right) \right\} + u_0. \tag{55}$$

Knowing that

$$\mathcal{L}^{-1} \left\{ \frac{1}{\sqrt{s}} \exp\left(-\frac{x}{\sqrt{D}} \sqrt{s}\right) \right\} = \frac{1}{\sqrt{\pi t}} \exp\left(-\frac{x^2}{4Dt}\right) \tag{56}$$

and

$$\mathcal{L}^{-1}\{G(s)K(s)\} = \int_0^t g(\tau)k(t-\tau)d\tau, \tag{57}$$

the concentration as a function of space and time is given by

$$u(x,t) = u_0 + \frac{1}{F\sqrt{\pi D}} \int_0^t j(\tau) \frac{1}{\sqrt{t-\tau}} \exp\left(-\frac{x^2}{4D(t-\tau)}\right) d\tau. \tag{58}$$

The concentration at the cathode ($x = 0$) is given by

$$u(0,t) = u_0 + \frac{1}{F\sqrt{\pi D}} \int_0^t j(\tau) \frac{1}{\sqrt{t-\tau}} d\tau. \tag{59}$$

Let u_{\min} be the minimum concentration required to start deposition, and ξ be the time when the minimum concentration is reached. Therefore,

$$u_{\min} = u_0 + \frac{1}{F\sqrt{\pi D}} \int_0^\xi j(\tau) \frac{1}{\sqrt{\xi-\tau}} d\tau. \tag{60}$$

Let $K = \frac{1}{2}(u_{\min} - u_0)F\sqrt{\pi D}$ be a constant characterizing the deposition onset.

C.1.1. Constant current density

For a constant current density j , the concentration is defined as

$$u(0,t) = u_0 + \frac{2j_0}{F\sqrt{\pi D}} \sqrt{t} \tag{61}$$

and the electric charge is defined as

$$Q(t) = \int_0^t j dt = j_0 t. \tag{62}$$

At the deposition onset, this yields the well-known Sand's equation (25):

$$\frac{1}{2}(u_{\min} - u_0)F\sqrt{\pi D} = j_0 \sqrt{\xi}. \tag{63}$$

The minimum charge can be derived accordingly:

$$Q_{\min} = j_0 \xi = \frac{K^2}{j_0} \tag{64}$$

C.1.2. Linear voltage ramp

If the voltage ramp is a linear function of time. Before the deposition, due to the Ohm's law, the current density is a linear function of time $j = \beta t + j_0$, and the concentration is expressed as

$$u(0,t) = u_0 + \frac{1}{F\sqrt{\pi D}} \int_0^t j(\tau) \frac{1}{\sqrt{t-\tau}} d\tau \tag{65}$$

$$= u_0 + \frac{1}{F\sqrt{\pi D}} \left[2j_0 \sqrt{t} + 2 \int_0^t \frac{\partial j}{\partial \tau} \sqrt{t-\tau} d\tau \right] \tag{66}$$

$$= u_0 + \frac{1}{F\sqrt{\pi D}} \left[2j_0 \sqrt{t} + \frac{4}{3} \beta t^{3/2} \right]. \tag{67}$$

At the deposition onset, the induction time is the solution of the following equation

$$\beta \xi^{3/2} + \frac{3}{2} j_0 \xi^{1/2} = \frac{3}{2} K. \tag{68}$$

Assuming that the voltage at the anode is initially 0, we can derive the following expression for the induction time

$$\xi = \beta^{-2/3} \left[\frac{3}{2} K \right]^{2/3} \tag{69}$$

Finally, the minimum electric charge for the linear voltage ramp is

$$Q_{\min} = \left(\frac{81}{128\beta} \right)^{1/3} K^{4/3}. \tag{70}$$

D. Baseline/Inference-informed model comparisons

We include here in Figures 14–15, all comparisons between the baseline and inference-informed models. This includes current and resistance comparisons for all six experiments. The inference-informed model performs better than the baseline model in all experiments, but the notably greater improvement on predictions for CC experiments. However, thickness prediction is still inaccurate in many cases.

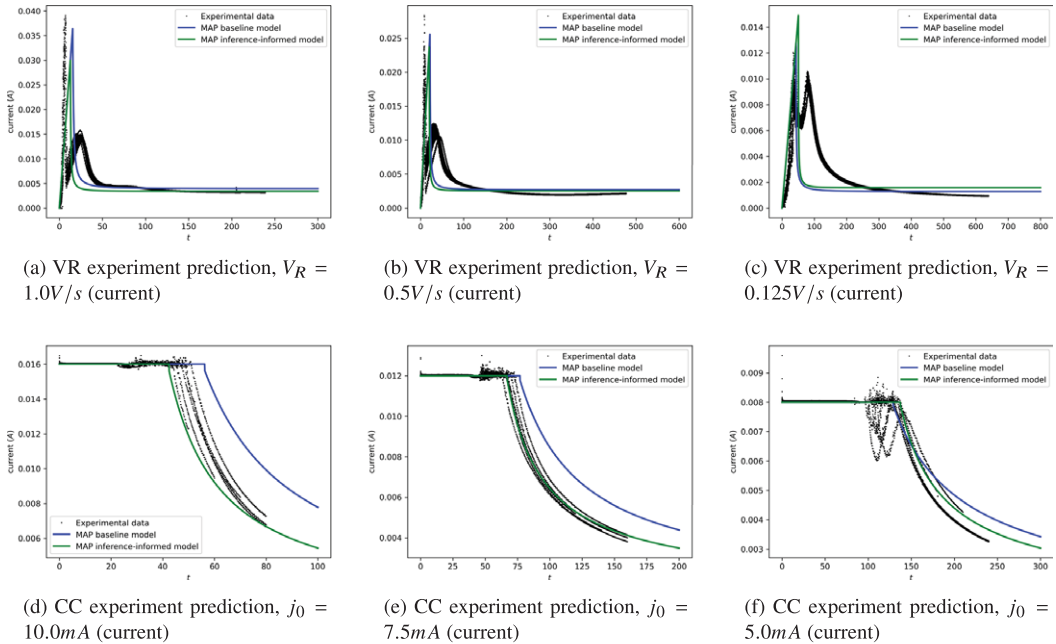


Figure 14. Comparisons between current prediction on the baseline model and inference-informed model at the MAP for each.

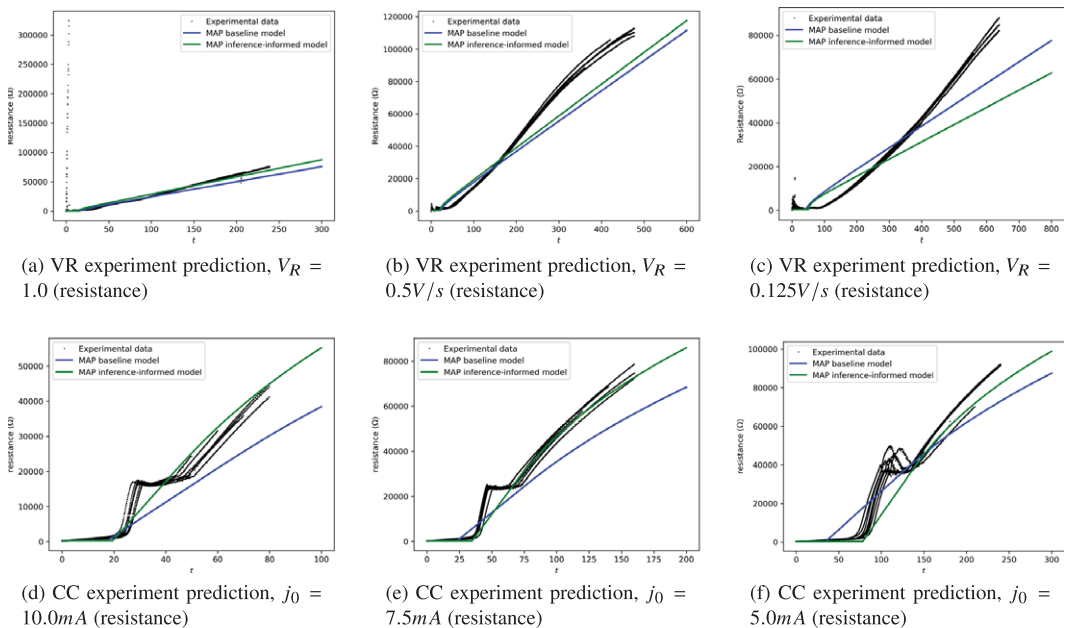


Figure 15. Comparisons between resistance prediction on the baseline model and inference-informed model at the MAP for each.

E. ML-augmented model with first peak

We present all prediction results on the ML-augmented model with first peak. Notably, predictions are poor for VR experiments in the low V_R regime but accurate for all other experimental configurations. In particular, thickness prediction is much more accurate compared to the baseline and inference-informed models. Predictions for current and resistance are included in Figures 16–17, but the ML model is trained with current data for the three voltage ramp experiments only.

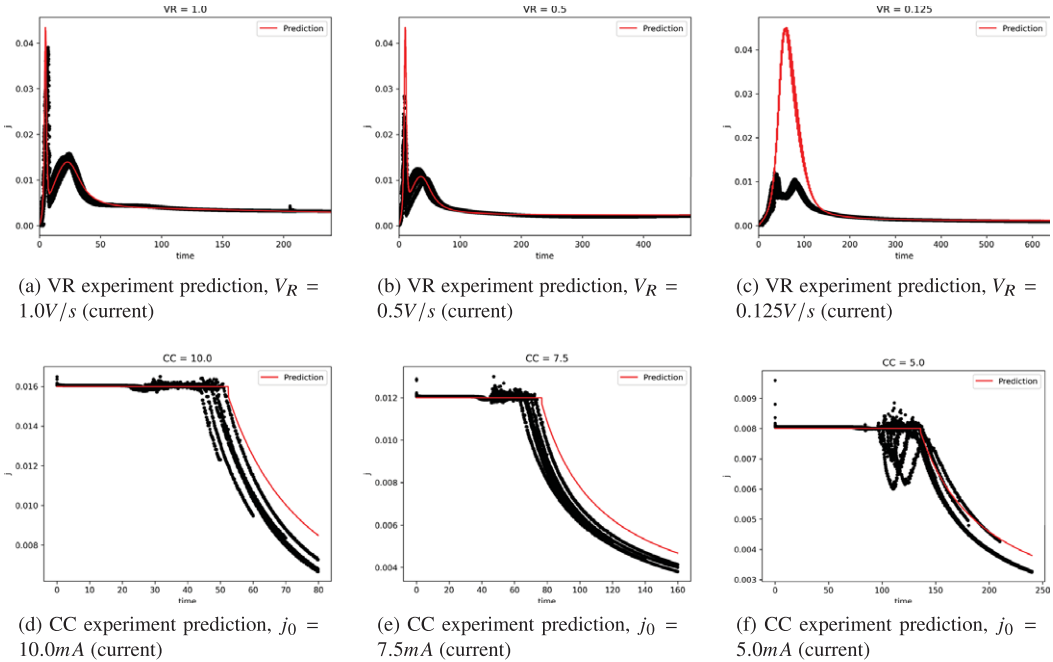


Figure 16. ML-augmented model with first peak current predictions compared to experimental data.

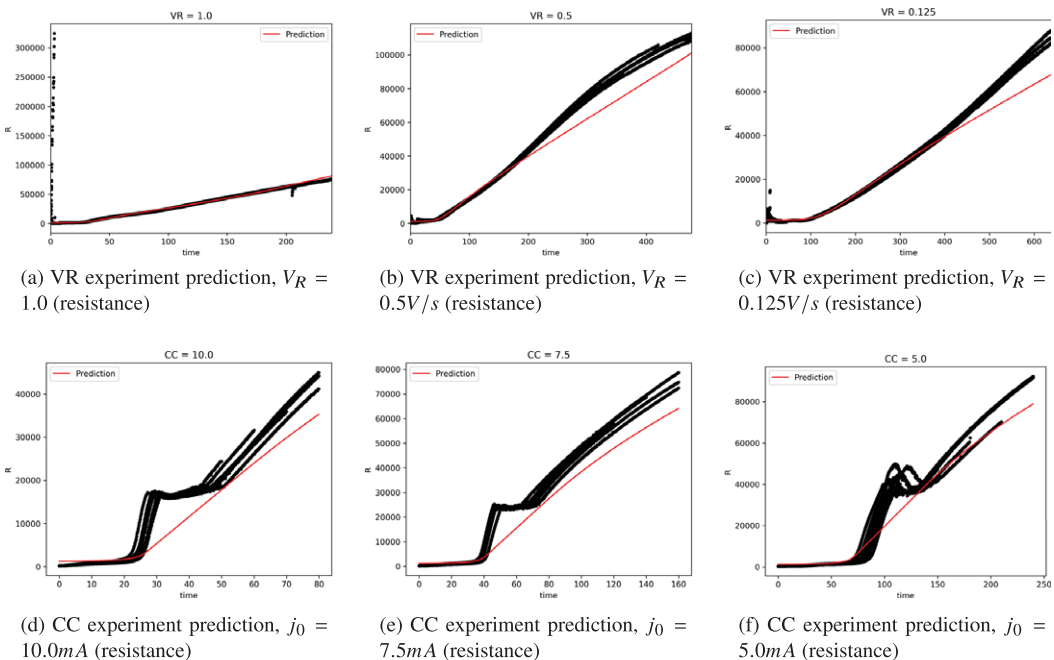


Figure 17. ML-augmented model with first peak resistance predictions compared to experimental data.

F. ML-augmented model without first peak

We present all prediction results on the ML-augmented model without the first peak model included. Removing the first peak model greatly improves model efficiency during prediction time (and training time). Predictions are still poor for VR experiments in the low V_R regime but improved over the baseline and inference-informed models. Again, thickness predictions are more accurate compared to previous models. Predictions for current, resistance, and thickness are included for all six experimental configurations in Figures 16–19, but the ML model is trained with current data for the 3 voltage ramp experiments only.

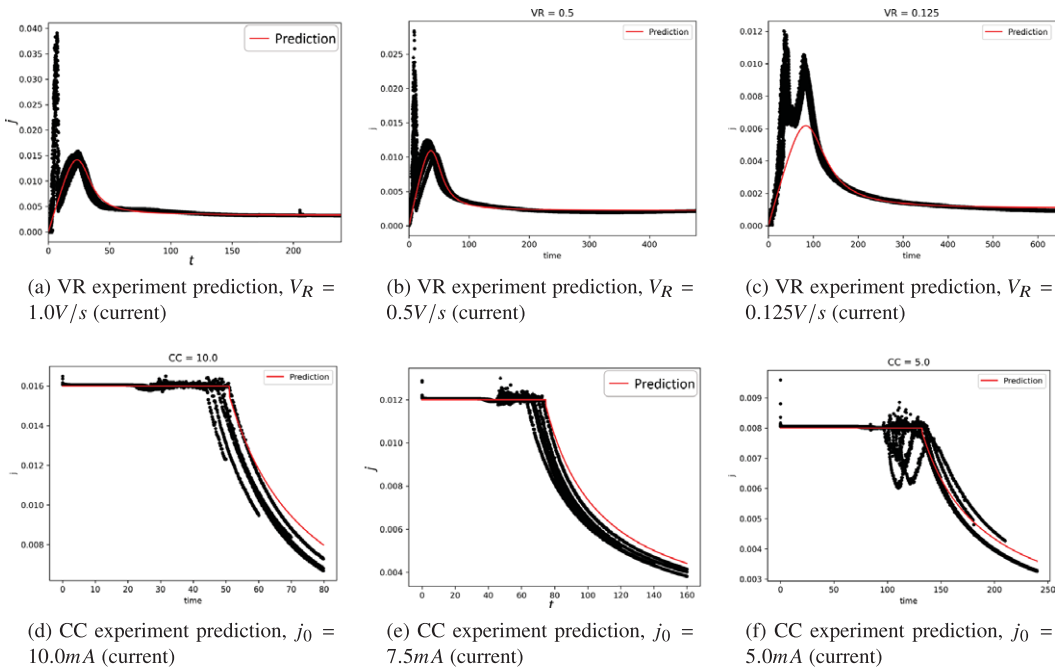


Figure 18. ML-augmented model without first peak current predictions compared to experimental data.

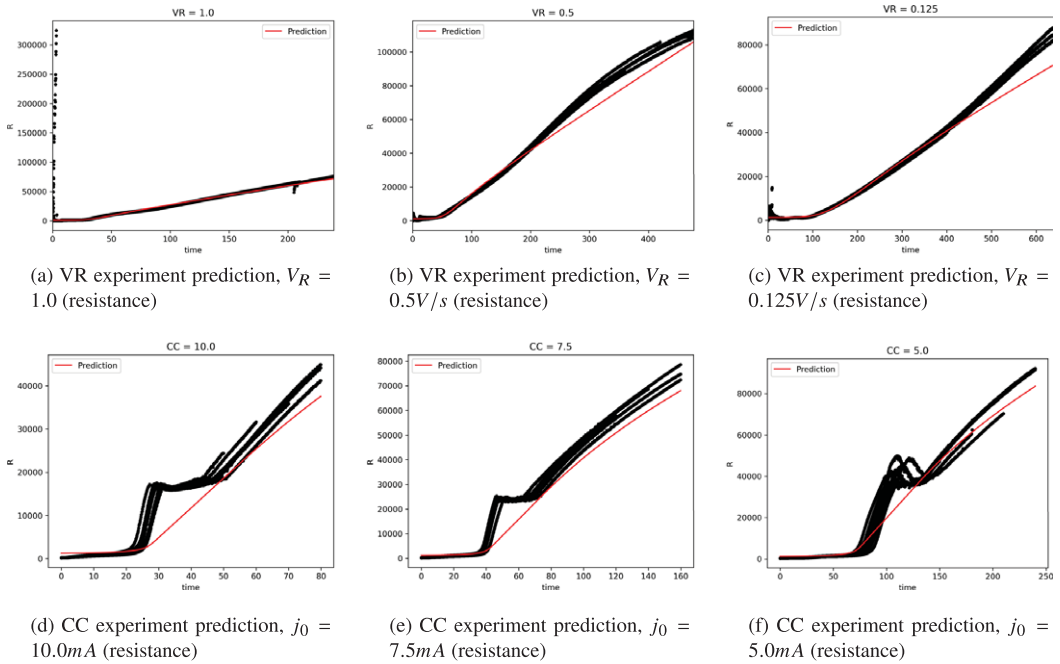


Figure 19. ML-augmented model without first peak resistance predictions compared to experimental data.