STARS

University of Central Florida
STARS

Electronic Theses and Dissertations, 2020-

2022

Lithium-ion Battery Prognosis with Variational Hybrid Physics-informed Neural Networks

Renato Giorgiani do Nascimento
University of Central Florida

Part of the Energy Systems Commons

Find similar works at: https://stars.library.ucf.edu/etd2020

University of Central Florida Libraries http://library.ucf.edu

This Doctoral Dissertation (Open Access) is brought to you for free and open access by STARS. It has been accepted for inclusion in Electronic Theses and Dissertations, 2020- by an authorized administrator of STARS. For more information, please contact STARS@ucf.edu.

STARS Citation

Giorgiani do Nascimento, Renato, "Lithium-ion Battery Prognosis with Variational Hybrid Physics-informed Neural Networks" (2022). *Electronic Theses and Dissertations, 2020-*. 1011.
https://stars.library.ucf.edu/etd2020/1011

Showcase of Text, Archives, Research & Scholarship

LITHIUM-ION BATTERY PROGNOSIS WITH VARIATIONAL HYBRID
PHYSICS-INFORMED NEURAL NETWORKS

by

RENATO G. NASCIMENTO
M.S. University of Central Florida, 2020
B.S. Universidade Estadual Paulista, 2015

A dissertation submitted in partial fulfilment of the requirements
for the degree of Doctor of Philosophy in Mechanical Engineering
in the Department of Mechanical and Aerospace Engineering
in the College of Engineering and Computer Science
at the University of Central Florida
Orlando, Florida

Spring Term
2022

Major Professor: Felipe A. C. Viana

ABSTRACT

Lithium-ion batteries are an increasingly popular source of power for many electric applications. Applications range from electric cars, driven by thousands of people every day, to existing and future air vehicles, such as unmanned aircraft vehicles (UAVs) and urban air mobility (UAM) drones. Therefore, robust modeling approaches are essential to ensure high reliability levels by monitoring battery state-of-charge (SOC) and forecasting the remaining useful life (RUL). Building principled-based models is challenging due to the complex electrochemistry that governs battery operation, which would entail computationally expensive models not suited for prognosis and health management applications. Alternatively, reduced-order models can be used and have the advantage of capturing the overall behavior of battery discharge, although they suffer from simplifications and residual discrepancy. We propose a hybrid solution for Li-ion battery discharge and aging prediction that directly implements models based on first-principle within modern recurrent neural networks. While reduced-order models describe part of the voltage discharge under constant or variable loading conditions, data-driven kernels reduce the gap between predictions and observations. We developed and validated our approach using the NASA Prognostics Data Repository Battery dataset, which contains experimental discharge data on Li-ion batteries obtained in a controlled environment. Our hybrid model tracks aging parameters connected to the residual capacity of the battery. In addition, we use a Bayesian approach to merge fleet-wide data in the form of priors with battery-specific discharge cycles, where the battery capacity is fully available (complete data) or only partially available (censored data). The model's predictive capability is monitored throughout battery usage. This way, our proposed approach indicates when significant updates to the hybrid model are needed. Our Bayesian implementation of the hybrid variational physics-informed neural network can reliably predict the battery's future residual capacity, even in cases where previous battery usage history is unknown.

To my wife Larissa and my daughter Sara.

ACKNOWLEDGMENTS

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1: INTRODUCTION

Motivation

Electric and hybrid systems are at the foundation of the current transportation transformation, in which small and large vehicles are expected to use electric or hybrid power-trains as a source of propulsion. This includes both ground vehicles, such as hybrid and electric cars driven by thousands of people every day, and future air vehicles [3–5]. Electric power-trains start with the energy storage unit, which most of the time is composed of Lithium-ion (Li-ion) or Lithium-polymer (Li-Po) battery packs. Although new concepts for energy storage are under consideration, including Lithium-air batteries [6] and nanotechnology applied to energy storage [7, 8]; Li-ion and Li-Po batteries remain the most common and consolidated way to provide energy to electric power-trains.

To achieve aircraft propulsion electrification, experts suggest a power-to-weight ratio of about 300 kW/kg by 2021 for 3-4 passenger Urban Air Mobility (UAM) type vehicles [9], which is expected to be achieved over the next years of technological development according to some qualitative estimates [9, 10]. Power over weight is not, however, the only requirement for a successful application of electric energy to aviation. Reliable tools to monitor the battery in real-time and perform diagnosis and prognosis are also equally important, for the obvious reason that aircraft loss of propulsion has more serious consequences than that of a ground vehicle.

(a) Thousands of people drive hybrid and electric cars every day.

(b) Batteries used to storage energy for electric power grids.

(c) Electric Urban Air Mobility vehicles fleets are expected over the next years.

Figure 1.1: Lithium-ion batteries are the most common and consolidated way to store energy in many applications.

Few Words about Prognosis: Data, Probabilistic Models, and First Principles

Machine learning has revolutionized the field of prognosis and health management [11, 12]. This is specially true in applications involving Lithium-ion batteries, where machine learning has been a key technology used in design [13–15], manufacturing and production optimization [16–18],
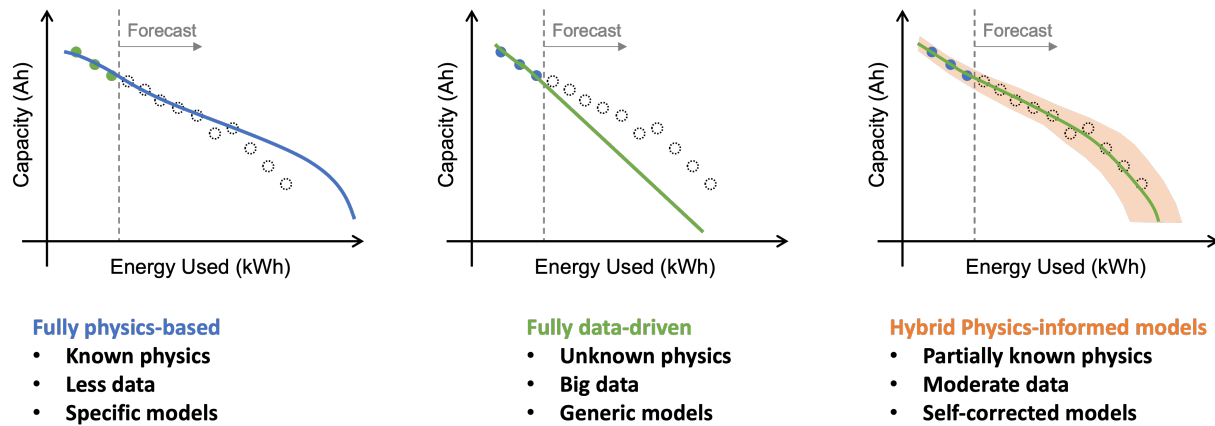
Figure 1.2: Illustration of different approaches for battery prognostics.

state-of-health estimation and remaining useful life prediction [19–21], and beyond. Machine learning can benefit from rich data coming from bench-top experiments and dedicated sensors placed in fielded products, compensating for potential lack of understanding of how systems work.

However, as the tasks become more complex, so do the machine learning solutions required to accomplish them. The cost of such solutions is the frequent need of large amount of well-curated data to guarantee good model generalization. To make matters worse, in system reliability, diagnostics, prognostics and health management; data can be incomplete, unbalanced, and present high uncertainty levels. Therefore, the use of machine learning can be very challenging.

Figure 1.2 illustrates different approaches to model the battery capacity of a specific battery with just a few observations. Employing a reduced order physics model, we can capture how an ideal battery would age. However, in reality, uncertainties from usage and battery properties might lead to discrepancies between the model and observed data. A pure data-driven method alone, would not be able to properly extrapolate and potentially offer unfeasible solutions. In this research, we propose a hybrid solution, in the form of a physics-informed neural network. In which a physics-based kernel guides the overall solution, and a data-driven part reduces the gap between the model

3

and the observed data.

Hybrid methodologies combining both worlds are now emerging in many fields, given recent advancements in software architectures to easily train data-driven algorithms and deep learning models, and the challenge in collecting useful data for critical systems (even with the advancement of the big data era, useful engineering data in many application areas are still small). The work presented in [22] goes in that direction, by extracting features related to the physics of charge of the battery and combining them with Gaussian process regression. Another example of hybrid approach is proposed in [23]. However, those examples show the implementation of a hybrid framework where data-driven and physics-based models are used in synergy, but they remain distinct in their definitions.

As a matter of fact, blending physics-based and data-driven techniques in a true sole model has its own challenges related to the identification of the merging point. Open research questions are "where does the flexibility of learning models outweighs the accuracy of physics?;" "can data-driven models compensate for the lack of knowledge of physical parameters?;" or "can physics guide and/or improve optimization of data-driven models (avoiding overfitting)?" Identifying parameters or a portion of a physics-based model that can be substituted by a data-driven algorithm is a challenging task. The latter can also cause over-fitting and reduce the extrapolation capability of the hybrid model. For that reason, careful initialization of the model parameters and constraints resembling physics in the estimation algorithm are paramount to ensure good fitting, but not obvious to define at times (as a matter of fact, the careful initialization of the parameters is also important in deep learning models for reasons related to gradient-based optimization and saturation of activation functions).

However, in real-time applications, learning models can compensate the natural inter-specimen variability typically observed in fleets of the same (or similar) systems, thus aiding uncertainty

4

quantification, and advances in deep learning training systems in the last decade made the adoption of large data-driven models very appealing, even though the risk of over-fitting in a limited-data scenario is high. Relying too heavily on deep learning models requires amount of data that are rarely available in the area of complex engineering systems and in the field of fault detection and failure prediction.

Moreover, deep learning models obfuscate the understanding of the model itself, reducing its interpretability. This might also hinder the model adoption in the field, as users may be reluctant to adopt a model they cannot interpret, as suggested by the latest research initiatives towards explainable and interpretable artificial intelligence in engineering [24, 25] and other fields [26, 27].

Scope and Organization of this Research

The research, conducted in this Doctoral Dissertation, has the following main objectives:

- Develop a hybrid physics-informed neural networks model for monitoring the state of charge (SOC) of Lithium-ion batteries.

- Implement and demonstrate that the proposed framework can accurately estimate voltage discharge cycles, track and forecast the aging of the battery over time.

- Illustrate how the resultant prognosis model can estimate the remaining useful capacity of lithium-ion batteries, and is able to forecast future discharge cycles, even with unbalanced or incomplete data, leveraging fleet-wide information.

The organization of this work is as follows. Chapter 2 presents a literature review of methods used for prognostics and health management of lithium-ion batteries. It includes the current developments in physics-informed neural networks and our hybrid implementation's novel aspects.

Chapter 3 discusses the implementation of the hybrid physics-informed model in a neural network graph. It introduces the reduced-order electrochemical kernel that governs battery operation and how a data-driven kernel accounts for the non-ideal voltage drop observed in experiments. Chapter 4 investigates how to track the aging of battery parameters with variational models. It shows how model ensembles can capture the variability between similar batteries and quantify the uncertainty estimates of future discharge cycles. Chapter 5 examines a probabilistic approach to capture the uncertainty in battery aging with regular usage data from random discharge cycles, which does not observe the actual battery capacity. Chapter 6 tackles the scenarios in which the battery usage history is incomplete or unknown. It explores how fleet-wide data of similar batteries and load profiles can be leveraged to estimate battery remaining useful capacity. Finally, Chapter 7 highlights the present research work's major conclusions and portrays the future work scope on this topic.

CHAPTER 2: LITERATURE REVIEW

Battery Prognosis

A number of recent research efforts on battery management systems included works on diagnosis and prognosis to monitor the state of charge (SOC) of the battery and predict the remaining time to discharge, or the aging associated with multiple charge-discharge cycles. Models based on first principles, empirical equations, and data have been tested in multiple scenarios, with their advantages and drawbacks. The need of knowledge of the physical phenomena at some degree of granularity from one side, and the need of data on the other side. Targeting this trade-off, hybrid approaches focus on combining physics-based methods with experimental data. Table 2.1 presents few examples of published papers found in literature reviewed in each approach.

Table 2.1: Approaches to build prognosis models to monitor battery state of charge and aging

Physics-based methods	Hybrid methods	Data-driven methods
Dynamics of li-ion cells with electro-chemical models, **Plett** [28]	System RUL with data-driven and model-based methods, **Liao et al.** [23]	Remaining useful life (RUL) with "mixed" neural networks, **Ma et al.** [29]
Battery equivalent circuit-based model and probabilistic estimation, **Saha et al.** [19]	Physics of charge combined with Gaussian process (GP), **Li et al.** [22]	Autoencoder model with large cloud data, **Li et al.** [30]
Battery degradation with mechanistic model, **Dubarry et al.** [31]		Surrogate modeling of capacity degradation with GP, **Valladares et al.** [32]

Plett [28] studied multiple mathematical models to capture the dynamics of lithium-ion battery cells. The author incorporated the Shepherd model, the Unnewehr universal model, and the Nernst

model, which combined describes the dynamic contributions due to open-circuit voltage, ohmic loss, polarization time constants, electro-chemical hysteresis, and the effects of temperature. He described the model in discrete-time state-space form and used an extended Kalman filtering (EKF) method to estimate the model parameters.

Saha et al. [19] employed a simple equivalent circuit-based model, estimating values for capacitance, impedance, and internal resistances. They used experimental data to fit the probabilistic kernels of a relevance vector machine algorithm. By embedding the battery internal states estimates in a particle filter system, they approximated the battery state-of-charge (SOC) and derived remaining-useful-life (RUL) estimates.

Dubarry et al. [31] investigated a mechanistic model analyzing different degradation modes scenarios. They developed a physicochemical model based on electrode degradation mode, cell design, and operating condition. By splitting their model into two layers, a top layer with cell configuration and degradation modes and a sub-layer describing the electrode behavior, they simulated multiple conditions and cell degradation modes.

Tang et al. [33] explored tracking the battery aging, the state-of-health (SOH), with an empirical approximation in the form of an exponential sum. The authors employed Bayesian Monte Carlo (BMC) method to identify the model parameters. They validated their approach with experimental data and showed good performance when using just a few discharge cycles for training.

Advanced machine learning and deep learning methods can achieve remarkable results in battery modeling, considering that a reasonable quantity of data is available. Li et al. [30] used a deep learning algorithm with an autoencoder model to prognosticate battery aging. The authors estimated the aging trajectories, borrowing the concept of rain-flow cycle counting from the fatigue of materials. Data requirements to train the autoencoder were satisfied by a cloud-based resource collecting data from a city electric bus fleet.

Valladares et al. [32] examined the use of surrogate modeling, applying Gaussian Process (GP) regression models to forecast the capacity degradation of battery cells and predict the battery initial specific capacity as a function of the electrode (cathode) active material. The authors also analyzed the role of multi-fidelity data with co-kriging surrogates, looking into enhancing battery prognostics with cheap data.

A pure data-driven approach was proposed by Shen et al. [34], where a convolutional neural network was used for capacity estimation of Li-ion batteries. Authors used the concept of ensemble learning and transfer learning to mitigate the lack of data to train their deep learning model. Other data-driven methods were proposed in [35], leveraging the concept of ensemble prognostics, and [29], where the term "hybrid neural network" refers to the fact that the network has both convolutional and long-short term memory layers, but has no physics in the model.

Hybrid approaches studies have shown that combining phenomenological knowledge with advanced data-driven methods can provide significant results without requiring large datasets. Li et al. [22] combined partial incremental capacity and Gaussian process regression (GPR) to forecast battery state-of-health. The authors extracted features related to the physics of charge of the battery with signal processing techniques. They applied GPR to model the battery capacity estimation using the features variables.

Liao et al. [23] explored a fusion of data-driven and model-based methods in a hybrid prognostics framework. The authors used an equivalent circuit-based model tracking battery degradation with the internal parameters impedance, capacitance, and internal resistances. Then they used data-driven methods to infer the unobservable measurements (internal system states), feeding them back to the model-based method to predict remaining useful life.

Despite advances reported in the literature, it is still difficult to build models that can predict battery end of discharge while accurately accounting for the effect of battery aging. Challenges include

effect of random-discharge cycles, effect of cumulative energy drawn from the battery, and inter-battery variability. On top of that, prognosis models ideally should come with uncertainty estimates that aid in decision making about how batteries are operated as well as the need for model updates.

Physics-informed Neural Networks

While physics-based models are built using the best understanding of a given phenomenon, limitations in terms of computational cost and even knowledge of physics often impose limitations to the predictive ability of such models [36, 37]. Under such limitations, model predictions do not always agree with observations, even after model parameters are calibrated [38]. Therefore, one can appreciate the importance of research in the field of uncertainty quantification. While most popular approaches lean towards statistical learning methods [39, 40], there has been growing interest in modern deep learning approaches, specially physics-informed neural networks [41].

As opposed to the popular use of neural networks, physics-informed neural networks operate on top of governing differential equations. Instead of only reducing the mismatch between neural network predictions and observed data, the training of physics-informed neural networks uses the laws of physics to help handling the reduced number of data points, and constrain the hyper-parameter space (see [42] for further reference). With incompressible fluids, for example, this is achieved by discarding non-realistic solutions violating the conservation of mass principle. Raissi [43] approximated the unknown of solution of partial differential equations by two deep neural networks. The first network acts as a prior on the unknown solution (enabling to avoid ill-conditioned and unstable numerical differentiation). The second network works as a fine approximation to the spatiotemporal solution. The methodology was tested on a variety of equations used in fluid mechanics, nonlinear acoustics, gas dynamics, and other fields.

Wu et al. [44] discussed, in depth, how to augment turbulence models with physics-informed machine learning. They demonstrated a procedure for computing mean flow features based on the integrity basis for mean flow tensors and propose using machine learning to predict the linear and nonlinear parts of the Reynolds stress tensor separately. They used the flow in a square duct and the flow over periodic hills to evaluate the performance of the proposed method.

Hesthaven and Ubbiali [45] proposed a non-intrusive reduced-basis method (using proper orthogonal decomposition and neural networks) for parametrized steady-state partial differential equations. The method extracts a reduced basis from a collection of snapshots through proper orthogonal decomposition and employs multilayer perceptrons to approximate the coefficients of the reduced model. They successfully tested the proposed method on the nonlinear Poisson equation in one and two spatial dimensions, and on two-dimensional cavity viscous flows, modeled through the steady incompressible Navier–Stokes equations.

Swischuk et al. [46] demonstrated through case studies (predictions of the flow around an airfoil and structural response of a composite panel) that proper orthogonal decomposition is an effective way to parametrize a high-dimensional output quantity of interest in order to define a low-dimensional map suitable for data-driven learning. They tested a variety of machine learning methods such as artificial neural networks, multivariate polynomial regression, k-nearest neighbor, and decision trees.

Samaniego et al. [47] presented an approach for discretization of partial differential equations using deep neural networks. Authors discuss similarities with collocation methods and how energy functionals found in many mechanics applications lend themselves into loss functions that can be used for training the deep neural networks. They presented numerical results for cases in linear elasticity, elastodynamics, hyperelasticity, phase field modeling of fracture, and piezoelectricity.

Goswami et al. [48] studied the usage of physics-informed neural networks to solve brittle fracture

problems. Authors propose the hybrid framework to reformulate the crack path prediction problem, where deep neural networks are guided through boundary conditions, and a loss function based on the variational energy. They showcase their approach on six different fracture mechanics problems, and proved the accuracy as well as the efficiency of the proposed method.

Research on physics-informed machine learning has been concentrated on solving differential equations and/or deriving them with deep neural networks. While these are important research thrusts, many engineering applications often impose limited and/or poor data sets and physics that is only partially understood. For example, in diagnosis and prognosis of industrial equipment (e.g., wind turbines, jet engines, airplanes, etc.), predictive models for estimating hardware distress are used to reduce unscheduled maintenance and downtime [49].

In this work, instead of using machine learning to replace the physics or fit the data, we will focus on how to explicitly code the physics-based domain knowledge into deep neural networks while using machine learning to quantify and compensate for discrepancy with respect to observations.

Proposed Contributions of This Research

This research addresses the issue of building a model that can be used for modeling the SOC of a battery within a mission as well as tracking and forecasting the battery useful life. We propose a hybrid architecture for battery prognostic by leveraging a physics-based model developed in [1, 50] and merging it with a data-driven model in a recurrent neural network (RNN) cell [51–53]. The result is a physics-informed trainable model for battery SOC that accounts for aging effects. The physics model is embedded in the RNN cell architecture, thus creating implicit constraints driven by physics to the cell. A portion of the original physics-based model is substituted by a multi-layer perceptron (MLP), which has the flexibility to learn from historical data. Then,

12

by leveraging a probabilistic framework based on variational inference [54], which enables the estimation of probability distributions of weights and biases of the RNN cell, the model provides a fleet-based ensemble prognosis, where data from different batteries of a fleet are used to extrapolate the reduction of the capacity of the battery that is being monitored.

One may identify similarities between our proposed hybrid physics-informed neural network approach and other methods that combine physics-based models and statistical learning. For example, Kalman filtering and its variants often use reduced-order models that are constantly updated as observed data becomes available such that changes in model parameters can be tracked. Within the battery literature, the interested reader is referred to [55–57].

Our approach differs in which we use the neural network machinery to allow for:

- explicit hybridization of the model, such that model-form uncertainty can be captured by data-driven layers;

- nonlinearities in the model which can be directly handled by backpropagation;

- tracking of model parameters through variational inference.

While the resulting model is fast to compute and can be used for predictions in real time, the training and update can be computationally intensive (few hours in today's computers). This shortcoming might not be there in some implementations of the Kalman filtering and we can see this as being the compromise due to the added capabilities of our hybrid approach.

CHAPTER 3: HANDLING MODEL-FORM UNCERTAINTY WITH HYBRID PHYSICS-INFORMED NEURAL NETWORK

In this chapter, we discuss the proposed method and the implementation of our hybrid solution for modeling Lithium ion batteries. As we discussed in the previous chapter, the battery operation is governed by very complex phenomena. In the context of prognostics and health monitoring, reduced-order models are generally applied to allow fast prediction applications. The simplifications in these models introduce deviations from actual battery operations and need to be accounted for to replicate experimental observations properly.

While advanced machine learning models can be used to model the battery behavior, the sole use of data-driven methods tends to require large datasets that are not always available. Moreover, these methods usually appear as "black box" models, lacking interpretability, which often is required in sensitive engineering applications.

We propose a hybrid solution in the form of a physics-informed neural network. A physics-based kernel guides the overall solution, and a data-driven part reduces the gap between the model and the observed data. We implement both the governing model equations, the Nernst's and Butler-Volmer equations, and the non-ideal voltage corrector factor, with a multilayer perceptron (MLP), as operations in the same neural network graph. This hybrid model, organized as repeating cells in a recurrent neural network (RNN), describes the discrete state-space representation of the battery operation, replicating discharge cycles time-series. We validate our approach against the experimental data provided in the NASA Randomized Battery Usage dataset [2].

Reduced-order Physics-based Model

The model adopted in this research was first proposed in [1] and then improved in [58]. As illustrated in Figure 3.1, this is an electrochemical model that estimates the battery output voltage as a sum of internal voltage terms, which considers the concentration of Lithium ions in each electrode. The purpose of the model was to work in real-time or quasi-real-time applications, and was designed with ordinary rather than partial differential equations, at the cost of some precision and accuracy loss. The use of ordinary differential equations facilitates the embedding of the model in fast prediction frameworks. Because of such simplifications, the model uses physics-based equations as well as some empirical derivations. As equations and descriptions are based on [1], the interested reader is referred to that article for more details and a thorough description of the model.

Figure 3.1: The electrochemical model assesses battery voltage terms, considering the Li-ion concentration in each electrode. Adopted from [1]

Equilibrium Potential

The model consists of the Nernst's equation for the equilibrium potential of a Li-ion battery:

$$V_{U,i} = U_0 + \frac{RT}{mF} \ln \frac{1 - x_i}{x_i} + V_{ni,i}\,. \tag{3.1}$$

where the electrode (negative or positive) is indicated by the subscript $i = \{n, p\}$; U_0 is the reference potential; R is the universal gas constant; T is the electrode temperature; m is the number of electrons transferred in the reaction; F is the Faraday constant; x is the mole fraction for the Lithium-intercalated host material; and $V_{ni,i}$ is the non-ideal voltage and activity correction term, null in ideal conditions. Details about V_{ni} will be provided later in the section. The mole fraction is computed as the ratio between the amount of Li-ion q in electrode i, and the amount of available (moving) Li-ions q^{max}:

$$x_i = q_i/q^{max} \qquad \text{and} \qquad q^{max} = q_n + q_p\,. \tag{3.2}$$

In order to accommodate the concentration gradient at the surface of the electrode, the total volume of the battery is split into two control volumes, bulk and surface, and the concentrations of Li-ions are calculated accordingly:

$$c_{b,i} = \frac{q_{b,i}}{v_{b,i}}\,, \quad c_{s,i} = \frac{q_{s,i}}{v_{s,i}}\,, \quad q_p = q_{s,p} + q_{b,p}\,, \quad q_n = q_{s,n} + q_{b,n}\,, \text{ and } \quad q^{max} = q_{s,p} + q_{b,p} + q_{s,n} + q_{b,n}\,. \tag{3.3}$$

The diffusion rate from the bulk to the surface is:

$$\dot{q}_{b\,s,i} = \frac{1}{D}\left(c_{b,i} - c_{s,i}\right), \tag{3.4}$$

where D is the diffusion constant. The rates of change of the charges q are calculated using the diffusion rate and the applied electric current i_{app}:

$$\dot{q}_{s,p} = i_{app} + \dot{q}_{b\,s,p}\,, \quad \dot{q}_{b,p} = -\dot{q}_{b\,s,p} + i_{app} - i_{app}\,, \quad \dot{q}_{s,n} = -i_{app} + \dot{q}_{b\,s,n}\,, \quad \dot{q}_{b,n} = -\dot{q}_{b\,s,n} + i_{app} - i_{app}\,.$$

$$(3.5)$$

The concentration overpotential is calculated using the Nernst's equation for the surface, which is by substituting x_i with $x_{s,i}$:

$$x_{s,i} = \frac{q_{s,i}}{q_{s,i}^{max}} \qquad \text{and} \qquad q_{s,i}^{max} = q^{max}\,\frac{v_{s,i}}{v_i}\,. \tag{3.6}$$

The solid-phase Ohmic resistance, electrolyte Ohmic resistance, and current collector resistance can be lumped together into R_0 to calculate the voltage drop: $V_0 = i_{app}\,R_0$, where i_{app} is the required current [1].

The surface overpotential is described by the Butler-Volmer equation [1], and the voltage $V_{\eta,i}$ can be expressed as:

$$V_{\eta,i} = \frac{RT}{F\alpha}\,\text{arcsinh}\left(\frac{J_i}{2J_{i0}}\right)\,. \tag{3.7}$$

This way, the battery output voltage is defined by:

$$V = V_{U,p} - V_{U,n} - V_0 - V_{\eta,p} - V_{\eta,n}\,, \tag{3.8}$$

which will serve as output of the physics-informed model.

Non-ideal Voltage, Internal Resistance, and Amount of Available Li-ions

As part of the simplifications used to obtain the described reduced-order physics-based model, the non-ideal voltage, battery resistance, and amount of available Li-ions are elements of the model fitted using experimental data. As already mentioned, the non-ideal voltage V_{ni} is the activity correction term needed to compute $V_{U,i}$. The model in [50, 1] fits the experimental data using the Redlich-Kister expansion

$$V_{ni,i}(x_i;\ \mathbf{A}_i) = \frac{1}{m\,F} \sum_{k=0}^{N_i} A_{k,i} \left((2x_i - 1)^{k+1} - \frac{2x_i k\,(1 - x_i)}{(2x_i - 1)^{1-k}} \right). \tag{3.9}$$

The mole fraction x_i is the independent variable, the coefficients $A_{k,i}$ are identified through data-fitting, and the number of elements in the sum N_i is empirically-derived as well. For the batteries tested in this work, the reference papers [50, 1] kept $N_p = 12$ and $N_n = 0$, thus using a constant non-ideal voltage for the negative side of the electrode.

Lumped internal resistance R_0 and maximum number of available Li-ions q^{max} drastically affect the prediction performance of the model. Such influence appears at both the *inter-* and the *intra-*specimen level; those two parameters change as the battery ages. The loss in active material can be represented by a drop in q^{max}, while an increase in R_0 is representative of a constant Ohmic drop, which causes an increase in total resistance independent of the battery charge state [58]. Such effects should be accounted for when training on multiple discharge curves of the same battery type. Moreover, different samples of the same fleet of batteries also have slightly different R_0 and q^{max}, thus varying from one battery to another and contributing to inter-specimen variability. That adds to the variability that can be found in small changes or imperfections of the manufacturing process and raw materials.

Physics-Informed Neural Network Model

This section details our approach to create a hybrid model that embeds the physics-based reduced order models in a neural network. By design, our proposed hybrid model leverages the physics expressed by the reduced-order models within a recurrent neural network (RNN). The hybrid model predictions are expected to capture major trends and could agree with experimental data at least in parts. In order to close any remaining systematic gap, we propose hybridizing the reduced-order model with data-driven layers, such as a multi-layer perceptron. The role of the data-driven layers is to characterize the model-form uncertainty in the reduced-order physics-based models. For a survey about physics-informed neural networks, the reader is referred to [59]. Here, emphasis goes onto the estimation of the non-ideal voltage $V_{ni,i}$, which, in the original model, was estimated with Eq. (3.9), as well as the total resistance R_0 (ohms) and the maximum charge q^{max} (coulombs), which values are driven by the aging of the battery.

Battery Model as a Recurrent Neural Network

A RNN extends the traditional feed forward networks to model time-dependent responses [60–62], as shown in Fig. 3.2a. RNNs repeatedly apply transformations to given states in a sequence such that

$$\begin{bmatrix} \mathbf{y}_t & \mathbf{h}_t \end{bmatrix}^\top = f\left(\mathbf{u}_t, \mathbf{y}_{t-1}, \mathbf{h}_{t-1}\right), \qquad (3.10)$$

where the subscript t represents the time discretization, $\mathbf{y} \in \mathbb{R}^{n_y}$ are the observable states, $\mathbf{h} \in \mathbb{R}^{n_h}$ are the internal states, $\mathbf{u} \in \mathbb{R}^{n_u}$ are input variables, and $f(.)$ defines the transition between time steps (function of input variables and previous states). In the recurrent neural network terminology, different implementations of $f(.)$ are referred to as cells.

Since the surrogate physics model describes the discrete state-space representation of the govern-
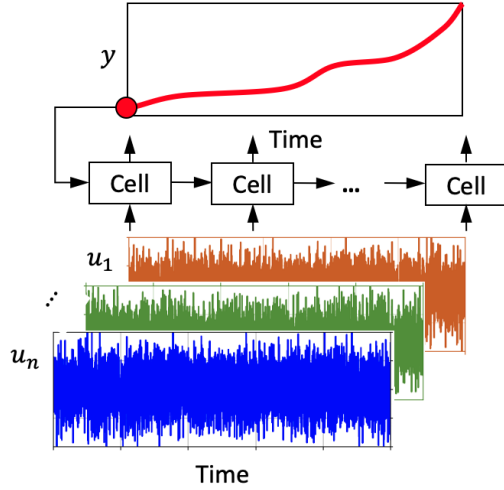
(a) RNNs cells repeatedly apply transformation to outputs, y, and inputs, $\mathbf{u}$.
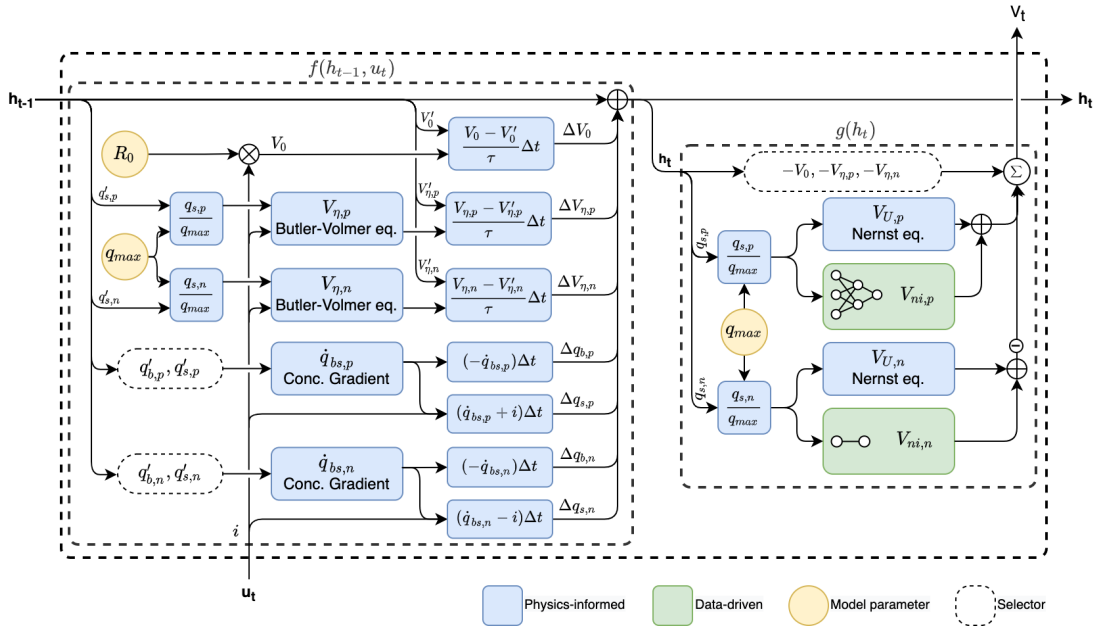
(b) Physics-informed RNN cell performing the state-space step-ahead prediction.

Figure 3.2: Physics-informed recurrent neural network.

ing set of ordinary differential equations, it allows the definition of a state update per time step. The approach follows the work presented in [51]; since RNNs naturally operate on discretized temporal sequences, they are suited to perform numerical integration. We proceed by designing a RNN that embeds the battery model, as illustrated in Fig. 3.2b. This design corresponds to the implementation of the physics-based model previously described as a directed graph model [63, Ch. 2] coded in the form of a RNN. The RNN cell task is to take the state vector at the previous time step, $\boldsymbol{h}_{t-1}$ and the input at the current time step, $\boldsymbol{u}_t$, and update the state vector $\boldsymbol{h}_t$ and the output y_t. The RNN cell produces a one-step-ahead prediction of the entire state-space model

$$\boldsymbol{h}_t = g_1(u_t, \boldsymbol{h}_{t-1}) \quad \text{and} \quad y_t = g_2(\boldsymbol{h}_t), \tag{3.11}$$

where the state, input, and output vectors are

$$\boldsymbol{h} = [T, V_o, V_{\eta,n}, V_{\eta,p}, q_{b,n}, q_{s,n}, q_{b,p}, q_{s,p}]^\top, \qquad u = i_{app}, \qquad \text{and} \qquad y = V. \tag{3.12}$$

In Fig. 3.2b, the cell receives the state vector at the previous time step ($\boldsymbol{h}_{t-1}$) and the input value at the current time step (u_t), and returns the state vector and the output voltage ($\boldsymbol{h}_t$ and V_t). The blue boxes are pure physics-blocks, which perform the same calculations of the physics-based model. The changes in the internal voltage states (ΔV_o, $\Delta V_{\eta,p}$, $\Delta V_{\eta,n}$) are obtained using an empirical time constant τ [1]. The dashed-rounded white boxes show variable selectors. The green boxes are representative of the two data-driven models used to estimate the non-ideal voltages, while the yellow circles represents the adjustable physical model parameters, q^{max} and R_0. The summation at the top of the cell is representative of Eq. (3.8).

21

The physics-based model previously discussed has two parameters and two data-driven components that need to be empirically adjusted based on observed data. These are the internal resistance R_0, the amount of available Li-ions q^{max}, and the non-ideal voltages $V_{ni,p}$ and $V_{ni,n}$, respectively. For example, in [58], authors used test results from specific reference discharge cycles to estimate R_0 and q^{max} as model parameters, as well as the coefficients of the Redlich-Kister expansion for $V_{ni,p}$ – i.e., $A_{k,i}$ in Eq. (3.9) – and coefficients of a linear regression model for $V_{ni,n}$.

Therefore, R_0 and q^{max} are kept as hybrid model parameters. Multi-layer perceptrons (MLPs) take the role of the Redlich-Kister expansion and linear regression models to capture the dynamic of the non-ideal voltage

$$V_{ni,n} = \text{MLP}_n(x_n; \mathbf{w}_n, \mathbf{b}_n) \qquad \text{and} \qquad V_{ni,p} = \text{MLP}_p(x_p; \mathbf{w}_p, \mathbf{b}_p) \,, \qquad (3.13)$$

where $\mathbf{w}_n$, $\mathbf{b}_n$, $\mathbf{w}_p$, and $\mathbf{b}_p$, are the MLP parameters. Vectors $\mathbf{w}$ represent the collection of weights of the MLP layers, while vectors $\mathbf{b}$ represent the collection of biases. Subscripts p and n refer to positive and negative electrode, respectively. Following typical MLP architectures, each node of the networks has its own weight and bias. The MLP is not restricted by the equation form of the Redlich-Kister expansion, and so it allows more flexibility in learning battery dynamics that are not captured by the physics-model. This advantage comes with the need for careful initialization of the MLP parameters, analysis of the model training during development to ensure the MLP is learning what is supposed to (the variability of the non-ideal voltage due to the activity coefficients), and not compensating for uncertainty or variability of other model parameters.

With the introduction of the MLPs in Fig. 3.2b, the model becomes a hybrid RNN with nodes that are physics-informed (e.g., Nernst and Butler-Volmer equations) and nodes that are data-driven

Table 3.1: Configuration of multi-layer perceptrons used to approximate the behavior of the non-ideal voltage ("tanh" and "lin" are the hyperbolic tangent and linear activation functions, respectively).

Layer	$V_{ni,p}$ (positive side)		$V_{ni,n}$ (negative side)	
	# neurons	activation	# neurons	activation
#1	8	tanh	1	linear
#2	4	tanh		
#3	1	tanh		

(MLPs for non-ideal voltage).

Table 3.1 details the design of each MLP used in the hybrid model. Similarly to the pure physics model, a linear model is used for the non-ideal voltage of the negative side, which is a MLP with one layer, one hidden unit and no activation function. The non-ideal voltage curve for the positive side is more complex, and therefore a MLP with 3 layers (8 hidden units in the first layer and 4 in the second hidden layers) with hyperbolic tangents as activation functions is used. All hidden units and output units have biases. Both MLPs take the mole fractions x_n, x_p, as inputs, and return the corresponding non-ideal voltages. The larger MLP for the estimate of the positive-side non-ideal voltage has 57 trainable parameters, while the MLP for the negative-side non-ideal voltage has only two trainable parameters (slope and bias of the linear model).

It is important to highlight that the hybrid physics-informed neural network models the time-dependent discharge of the battery. As such, the input and output variables are the time histories of current and battery voltage. These are the only two observable quantities. All states used to describe the physics, including $\boldsymbol{h} = [T, V_o, V_{\eta,n}, V_{\eta,p}, q_{b,n}, q_{s,n}, q_{b,p}, q_{s,p}]^\top$, $V_{ni,p}$ and $V_{ni,n}$, are hidden variables. This makes the training of our hybrid physics-informed neural network very interesting, as $V_{ni,p}$ and $V_{ni,n}$ are never observed directly and q^{max} and R_0 are very deep in the graph (see Fig. 3.2b).

Results and Discussions

Our implementation is all done in TensorFlow[1] (version 2.3) using the Python application programming interface. The uncertainty quantification and probabilistic modeling is performed using the TensorFlow Probability library [2] (version 0.11.0). The next subsections will detail how the model is fitted and validated as well as its use for random discharge prediction. The fitted hybrid physics-informed neural network model was validated against the physics-based model, which is publicly available in the NASA's *Prognostics Model Library* [64].

Data Description

The Randomized Battery Usage dataset we utilized for validation contains data from battery discharge experiments in a controller environment (including constant environment temperature) [2]. It has been widely utilized in the prognosis and health management domain, and therefore, we report only a brief summary hereafter.

All batteries compose of a single cell, with a maximum voltage of approximately 4.2 V when fully charged, and all tests stopped when the output reached 3.2 V. During the constant-loading tests, batteries were subject to 1 A current draw, while during the random-loading tests, the input current was randomized between 1 and 4 A, using a uniform distribution. Each sample was held constant for 300 seconds. So, $i = 1$ A during the constant-loading experiments and $i \sim \mathcal{U}[1, 4]$ A during the random-loading experiments. Figure 3.3a shows an example of full test history showing reference discharge cycles in green and random discharge cycles in orange. Figure 3.3b shows an example of constant (top) and random (bottom) loading discharge cycle.

---

[1]www.tensorflow.org

[2]www.tensorflow.org/probability
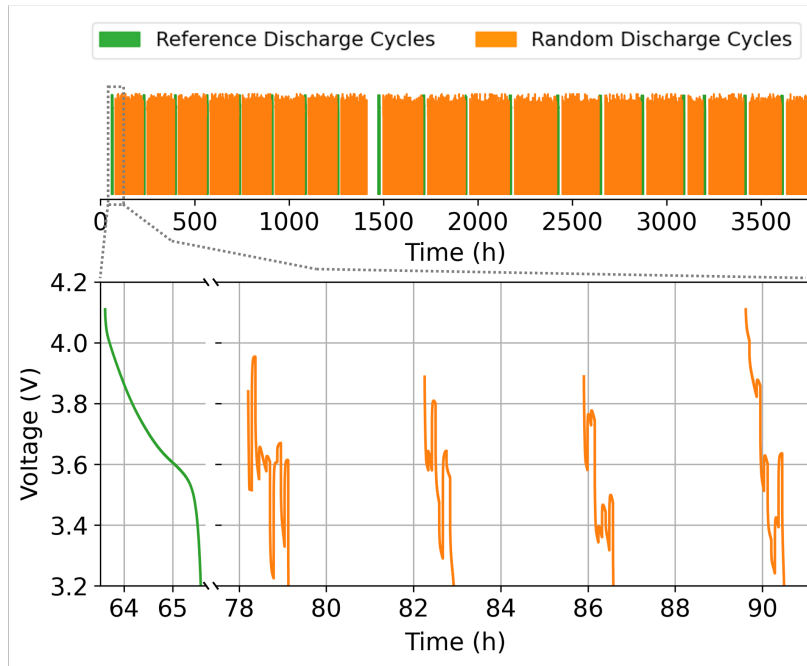
(a) Full history of discharge cycles for a single battery.
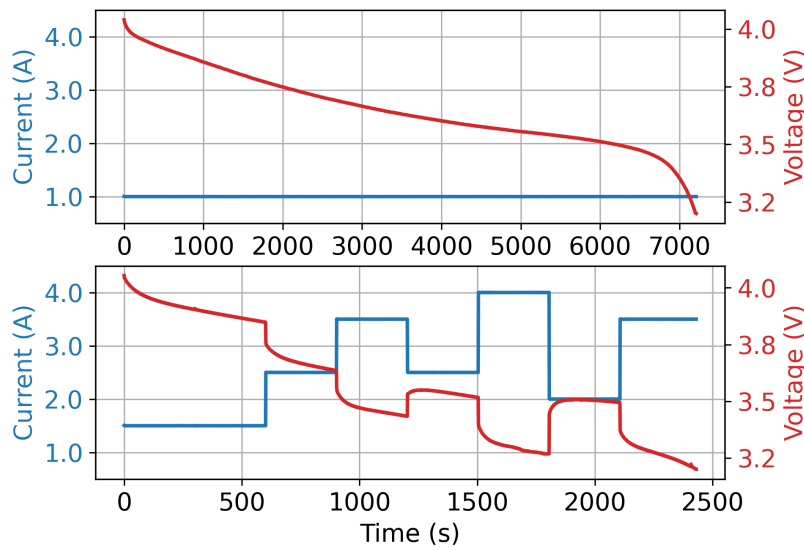
(b) Example of constant (top panel) and randomized (bottom panel) loading discharge cycles.

Figure 3.3: NASA Prognostics Data Repository - Randomized Battery Usage Data-set [2].

The first step consisted into training our hybrid physics-informed neural network, which simultaneously optimizes the values of q^{max}, R_0, and the MLP parameters ($\mathbf{w}_n$, $\mathbf{b}_n$, $\mathbf{w}_p$, and $\mathbf{b}_p$). We chose to do so using the constant-loading data [65]. The first 3 constant-loading discharge curves are extracted from each of the 12 batteries in the data set [65]. Each curve was generated with a current draw of 1 A from the fully charged condition of 4.2 V down to a value of 3.2 V, when the tests stopped. Then, the mean squared error is used as loss function Λ:

$$\Lambda = \frac{1}{N} \left(\mathbf{V} - \hat{\mathbf{V}} \right)^T \left(\mathbf{V} - \hat{\mathbf{V}} \right), \tag{3.14}$$

where N is the number of observations, $\mathbf{V}$ are the observed battery voltages over time, and $\hat{\mathbf{V}}$ are the predicted battery voltages by the physics-informed neural network. The model is trained with the Adam optimizer [66] set with a learning rate of 5×10^{-3} and 3000 epochs. All other Adam parameters were kept equal to their default values.

The physics-based portion of the model is initialized with the same physical and empirical parameters suggested for the physics-based surrogate model by original paper [1]. Battery parameters R_0 and q^{max} were initialized to 1.5×10^{-1} $\mathbf{\Omega}$ and 1.4×10^4 $\mathbf{C}$, respectively, which are reasonable values for the specific type of battery. We introduced further prior knowledge of the system by initializing the MLP for the positive electrode as a piece-wise linear curve representative of the decreasing non-ideal voltage over time. Such an initialization constitute a fundamental step in the design of the hybrid framework. Thus, the local optima the MLP_p converges to generates a curve resembling the one obtained by fitting experimental data with more traditional procedures, like the Redlich-Kister expansion in (3.9). In our experience, a random initialization of the MLP for $V_{ni,p}$ can pose problems such as poor convergence of model parameter optimization, overfitting of data, and lacking of generalization. One might argue this is not a problem only for neural networks

(or physics-informed neural networks for that matter), but for nonlinear regression in general. It is important to keep the range of $V_{ni,p}$ between -1 and 1 (as MLPs are normalized). However, the exact value of $V_{ni,p}$ for $x_p = 0$ and $x_p = 1$ are not as important (as long as they are close to 1 and -1, respectively). Likewise, the exact location of the knee of the bilinear curve is not as important. Even these simple incorporation of physically guided initialization was enough to improve optimization of MLP_p parameters. Once the MLP_p is initialized, the training process using a few discharge cycle curves helps fine tuning both MLPs to capture the underlying dynamics of the decreasing non-ideal voltage. The MLP for the negative electrode is the simplest linear model (slope and bias), and so the default Tensorflow initialization parameters is used.

A total of 36 discharge curves were used to simultaneously optimize the values of the both MLP parameters as well as q^{max} and R_0. This is an important feature of our physics-informed neural network. The estimation of physical model parameters and the optimization of data-driven model parameters is performed simultaneously using the time histories directly. This poses an advantage in many engineering applications, as the need for large data sets is reduced and compensated by the use of physics-based models.

Given that q^{max} and R_0 are estimated based on experimental data, slight variation in estimated values due to sample-to-sample variation is expected, even if data comes from the same battery type. Under such an assumption, the pair (q^{max}, R_0) can be used as surrogate to learn and represent the inter-specimen variability observed in specimens of the same battery type. Thus, a value pair of (q^{max}, R_0) is representative of a particular battery and the values in the battery RNN cell are vectorized, creating a pair (q^{max}, R_0) for each discharge curve. By so doing, after training, a distribution of the pair (q^{max}, R_0) is obtained while the MLPs learns the fundamental behavior of the non-ideal voltage for the battery type. Therefore, the trained MLPs are representative of the battery fleet, regardless of the input time history.

(a) Constant-loading discharge samples curves used for training (top panel) and examples of model prediction (bottom panel). The solid curves underneath the dashed curves are the ground truth curves used for model performance assessment.

(b) End-of-discharge prediction over the 36 sample curves used for training (dashed gray line shows +/- 1% band)
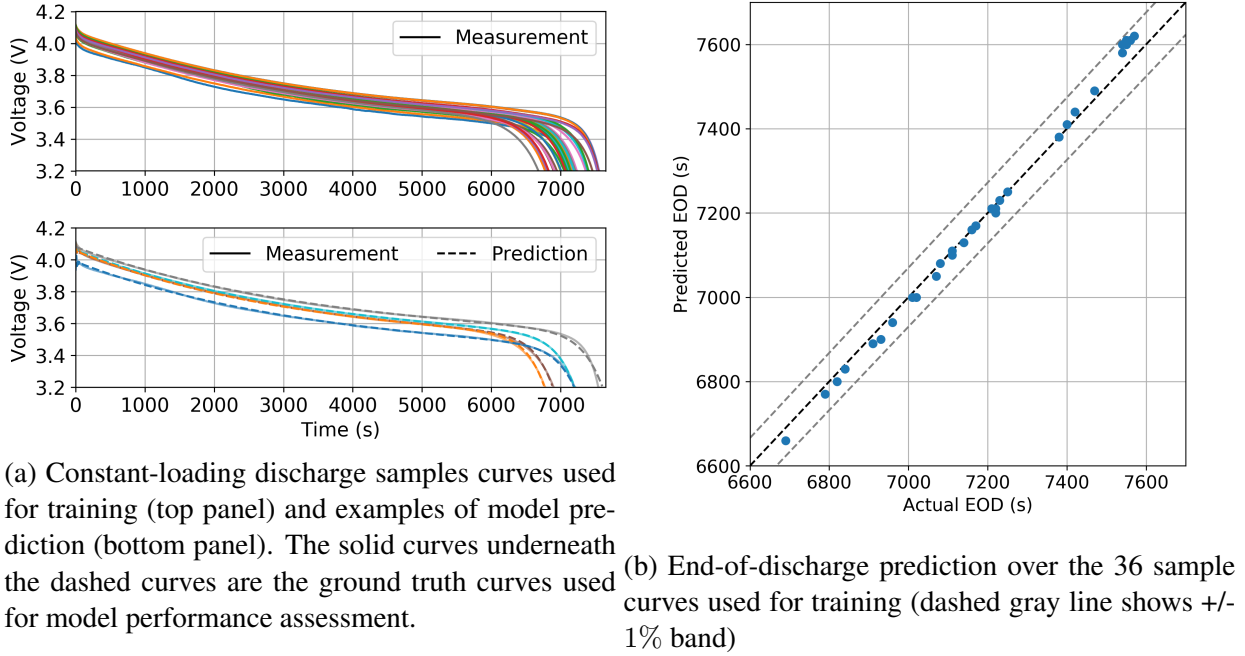
Figure 3.4: Model performance on constant-loading discharge samples.

First, we analyzed qualitatively the performance of the trained model in replicating constant-loading discharge curves. Figure 3.4a contrasts the observed data and model predictions. The solid curves underneath the dashed curves are the ground truth curves used for performance assessment. The root mean square error over the entire curve, averaged for all curves, is 6.7×10^{-3} V. The model accuracy in predicting the end of discharge can be assessed by looking at the time needed to reach a voltage of 3.2 V. As shown in Fig. 3.4b, the end-of-discharge predictions appear to be well-aligned with the true values, with a root mean square error of 27.5 seconds, in a range of 6700 to 7600 seconds.

We hypothesize that the non-ideal voltage models (MLP_n and MLP_p) can be shared across different batteries and do no change over time (i.e., no aging effect). In order to test this hypothesis, we carry out a cross validation study. In this process, we compare the distribution of the voltage out-

28

put for different training sets and assure the voltage output (from the sets left out for validation) is within the expected confidence bands obtained using the Wilson interval. We chose a k-fold cross validation approach to assess the robustness and generalization ability of the model over constant-loading discharge curves. The 36 samples at constant loading originally used for the first model fitting were split into 30 curves for training, and 6 for cross validation. The subsets (30-6) were shuffled to allow new combinations of training and validation curves at each fold. It is important to point out that in this k-fold cross validation the 6 curves left out are not used in the training of the model, which will only use the respective set of 30 curves. As this process is repeated and the subsets of (30-6) curves are shuffled, the robustness of the model with regards to the training data is assessed as well as its prediction capabilities outside the training set. Further details about cross validation can be found in [67, 68].

For reference, the model is trained using all 36 sample curves and obtained a MLP curve, one distribution for q^{max} and one for R_0. The dispersion of the MLPs' output is quantified when the model is trained with different folds. Each fold produces a value of q^{max} and one of R_0, and new parameters of the MLPs[3]. By comparing the MLP predictions for each fold, the robustness of the model against different combinations of curves in the training set is assessed. If that were not the case, then our approach of using one MLP model for the population build upon similar batteries would be invalid. The MLP_p outputs appear very similar to each other and close to the reference MLP_p curve, suggesting that using one MLP_p model for the population (multiple batteries with similar characteristics) can be considered a good approximation. The average root mean square error between the 6-fold-learned MLPs and the reference output curve is 1.22×10^{-3}, and the average standard deviation of the six curves is 8.94×10^{-3} – i.e., 0.45% of the non-ideal voltage

---

[3]Although both MLPs for positive and negative sides are trained, we show here only the output curves of the MLP for the positive side of the electrode.We did not encounter any challenge when training the linear model used for the negative side, MLP_n, and the variability observed in the MLP_n output does not influence the Battery model output nearly as much as the MLP used for the positive side of the electrode.

Table 3.2: Root mean square error over the validation set in the k-fold cross validation study.

k-fold [-]	Root mean square error $\times 10^{-3}$ [V]
1	7.5
2	7.7
3	6.9
4	7.3
5	7.9
6	7.9
All data in training	6.8

normalized range ($\approx$ 2V). Table 3.2 shows the mean square error between the predicted output voltage curve, and the observed output voltage curve, averaged over all the 6 validation curves.

We proceeded to further validate the output voltage predictions at constant loading by examining the probabilistic coverage of the predicted curves. The 30 sets of trainable elements (MLP parameters, q^{max}, R_0) obtained with the k-fold cross validation are then used to predict the confidence interval of the output voltage over time, thus producing a range of plausible voltage values at each step of the discharge cycle. The number of voltage data points of the remaining 6 validation samples that fall outside those confidence intervals is counted. Assuming a confidence intervals $1 - \alpha$, it is expected that the percentage of points outside the interval approximately equal to $\alpha \cdot 100\%$. The percentage points that should fall within the confidence interval is obtained using the Wilson interval [69, 70]:

$$\text{CI} = \hat{p} \pm z_{\alpha/2}\sqrt{\hat{p}(1 - \hat{p})/N}\,, \tag{3.15}$$

where $\hat{p}$ is the approximated target probability; $z_{\alpha/2}$ is the $\alpha/2$ critical value for the standard normal distribution, which denotes the $(1 - \alpha) \times 100^{th}$ percentile; and N is the number of observations.

In this case, since 6 curves out of 36 for cross validation are used, $N = 36$ and $\hat{p} = 100 \times 6/36 \approx$ 16.7%. Using the Wilson interval, $\hat{p}$ is expected to be in $16.7\% \pm 14.6\%$ interval (considering one

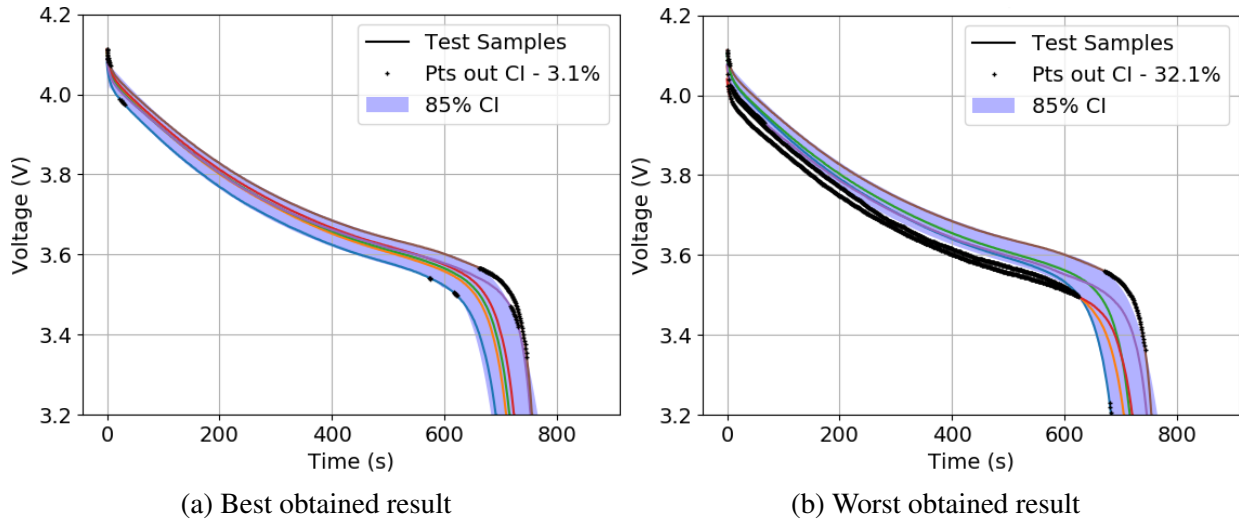(a) Best obtained result　　　　　　　　　　(b) Worst obtained result

Figure 3.5: Assessment of the model generalization by evaluating how many points of a experimental discharge curve fall outside the confidence intervals computed by the interrogation of the hybrid model with 30 different sets of model parameters. The two panels show results obtained with the data set subdivision of 30 reference curves and 6 validation curves.

standard deviation). Figure 3.5 shows two examples of this validation process for the best case (points outside the confidence interval are 3.1% of the entire curve) and the worst case (32.1% outside the CI) among the 6-fold validation. On average, 16.99% of the output voltage data points fall outside the CI, with a large standard deviation of 10.74%. Therefore, the Wilson interval still defines the result, $16.99\% \pm 10.74\%$, as acceptable. The apparently large standard deviation is just an artifact of the small sample size (6-curves) and it does not relate to our model predictions. The predictive capabilities of our hybrid model are confirmed by this cross-validation analysis.

Predictions of Random-Loading Discharge

The model over random-loading discharge curves is validated using 8 batteries from the same dataset. The parameter pair (q^{max}, R_0) comes from the reference discharge curve for the same

battery (at constant-loading), and it is used to estimate the voltage drops under random-loading. Therefore, the first random-loading test curves carried out right after the constant-loading tests are selected. By so doing, the effect of aging over the parameter pair (q^{max}, R_0) is minimized.

Predictions for all 8 batteries under random-loading conditions are evaluated individually. Figure 3.6 shows the prediction of the first random-loading discharge curves, after the batteries were tested in constant-loading conditions. The root mean square error between the observed voltage values (thick black line) and the predicted curve (dashed-black line) is always below 5.5×10^{-2} V and below 0.088% of the true voltage, except for battery $\# 1$, where larger discrepancies between model and real data were observed. We did not identify the reason for such a larger error in that particular battery, but we noticed that battery $\#1$ behaves slightly differently from the others and not adhering to the trends observed in the rest of the population.
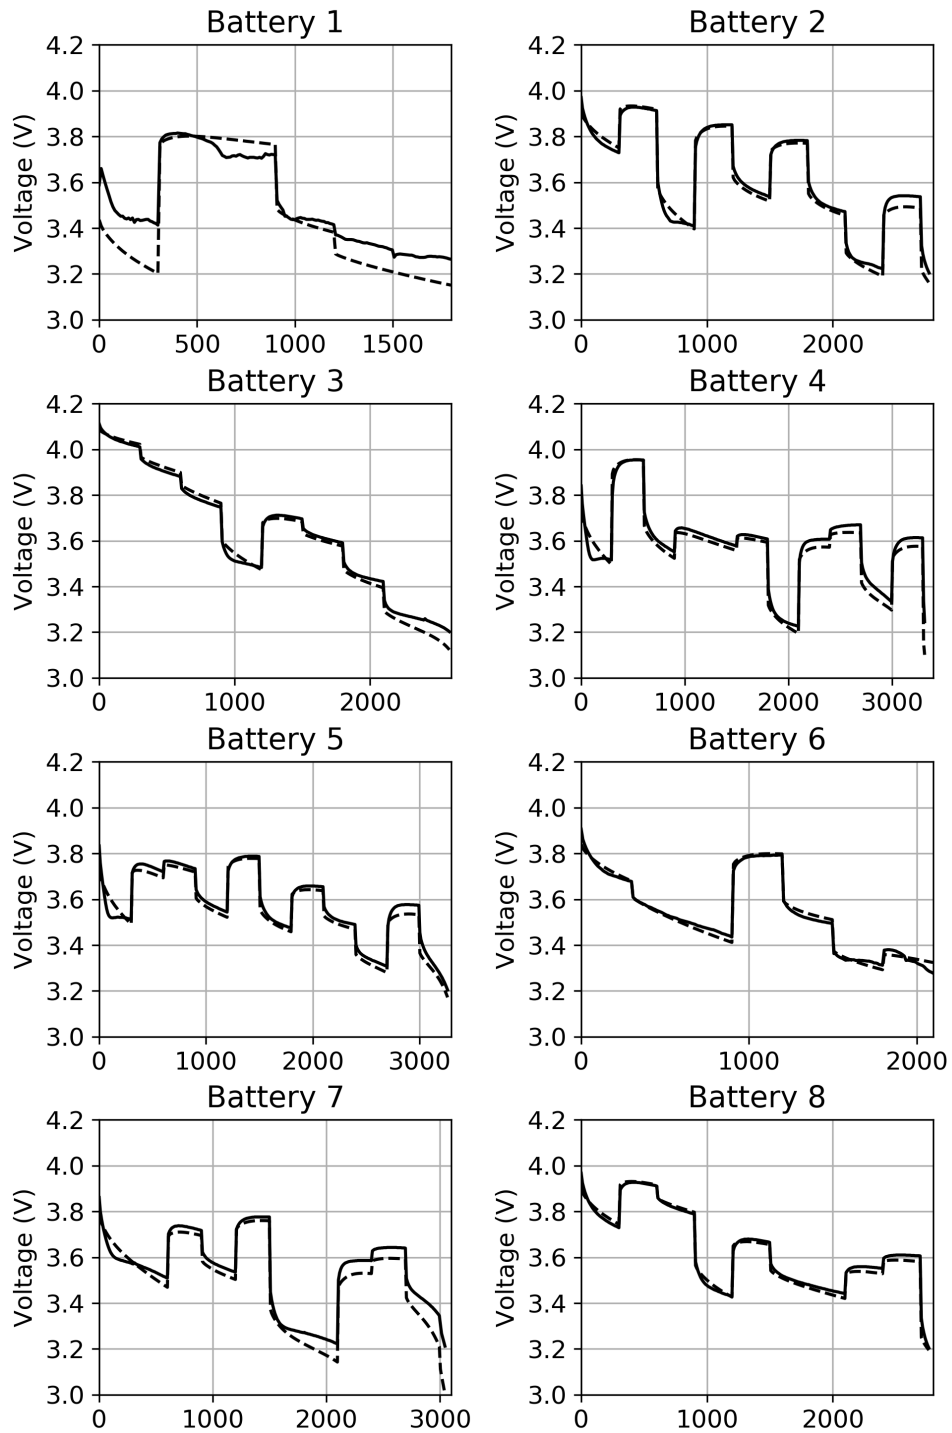
Figure 3.6: Prediction of first random-loading discharge curves for batteries 1 to 8 from model trained only with reference discharge curves (constant load).

CHAPTER 4: ACCOUNTING FOR BATTERY AGING WITH VARIATIONAL MODELS

The main motivation of this chapter is to overcome the limitations of model prediction as the batteries age. While we have observed that data-driven elements of our hybrid model, i.e. non-ideal internal voltage, do not vary over time; model parameters such as internal resistance and available ions tend to drift. Given the physics model structure, battery aging can be captured in part by q^{max} and R_0. Reductions in q^{max} accelerate the rate of discharge for the same current draw, and increases in resistance R_0 lead to drop in voltage [1]. As shown in Fig. 4.1, shortly after a reference discharge cycle (left panel), the values of q^{max} and R_0 are updated and predictions given by the model are in good agreement with observed data (middle panel). As the battery accumulates hours of operation, the model prediction starts disagreeing with the observed data (right panel).
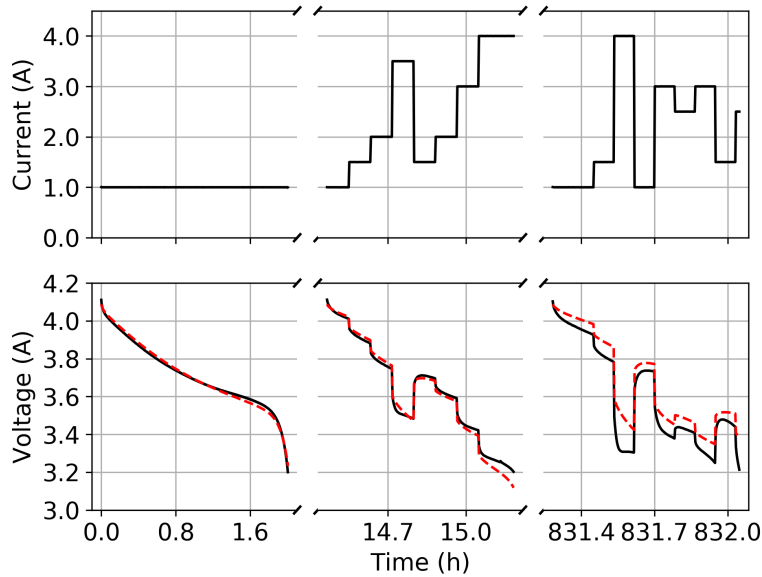

Figure 4.1: Prediction under constant (left panel) and random input current loading near the discharge cycle used to learn the model parameters (q^{max}, R_0) (middle panel) and further way in time (right panel). Black curves are observed data and red curves are predictions.

The aging effect does not limit the capability of the model trained with the proposed approach. In practice, we propose that one tracks the model prediction versus observed data. Then, the latest discharge curves can then be used to update the pair (q^{max}, R_0) by fitting the model, optimizing only the battery parameters pair while keeping the non-ideal voltage V_{ni} weights fixed. By doing so, the updated model predictions will account for the aging effect introduced by the latest loading curve. This continuous learning procedure is suitable for one-step-ahead prediction of aging. However, it does not enable far-ahead forecast of the battery aging, represented by a reduction in maximum capacity. We address how to forecast aging by using a fleet-based approach and ensemble modeling in the next section.

Aging Parameters and Cumulative Energy

Battery aging is manifested in the variations of the parameter pair (q^{max}, R_0) as loading cycles accumulate in the battery life. Therefore, we propose using the pair as a proxy indicator for aging. In terms of modeling, this implies that (q^{max}, R_0) are updated throughout the battery life, while the parameters for the two MLPs (which characterize interval voltage) are kept the same.

The pair (q^{max}, R_0) is updated using the sequential, constant-loading discharge curves available in the dataset for each battery. These reference-discharge curves are followed by a number of random-loading discharge, after which a new reference-discharge is performed. This information allowed us to characterize the changes in q^{max} and R_0 as the batteries accumulated loading cycles. Figure 4.2a shows that, as expected, q^{max} decreased over time (although not shown, as expected, R_0 increased over time) The rate of decay is visually similar for all the batteries since they have same characteristics and experienced similar loading conditions.

Although interesting, curves of q^{max} and R_0 as a function of time might not be as useful in real

life, as they do not capture the dependency with regards to required power per cycle. Instead, we suggest using cumulative energy as independent variable. With that, operators will be able to perform analysis and trade-off studies regarding mission dispatch and remaining useful life. Figure 4.2b shows battery capacity and q^{max} as a function of the cumulative energy. Circles refer to the capacity values, while full dots refer to q^{max} values, and their magnitude is shown in the double-y axis of the panel. The figure shows data for all batteries, emphasizing the similarities among them. The clear correlation can, in principle, allow us to model aging by modeling the decay of its two proxies (q^{max} and R_0) as a function of the past energy drawn. Figure 4.2b also suggests that up to 1 kWh, all batteries show a very narrow dispersion of the capacity and q^{max} drop, manifesting low inter-specimen variability. The curves then start diverging and the inter-specimen variability increases. Nevertheless, battery capacity and q^{max} trends remain highly correlated.

As previously shown, once q^{max} and R_0 are tracked, predictions for random discharge cycles are expected to improve. Figure 4.2c illustrates the adjustment in the same random discharge cycle shown in Fig 4.1 (right panel). The bias in the model prediction due to the misestimation of q^{max} and R_0 is removed once these parameters are updated.

Assuming that q^{max} and R_0 can be tracked for an existing set of batteries, we propose using an ensemble model that incorporates information from the set batteries in the prediction of the battery that is currently being monitored. The ensemble model combines the estimates from recorded q^{max} and R_0 curves with the data points available for the current operating battery, and produce the distribution *a-posteriori* of the future battery capacity. The underlying hypothesis is that the model has access to data from batteries (of the same type) that were already being used and therefore they are in a later aging state. Early in the history of the battery usage, this approach has the advantage of leveraging the observed behavior past 1 kWh from other batteries, and therefore capturing nonlinearities. Details of the developed ensemble is discussed next.
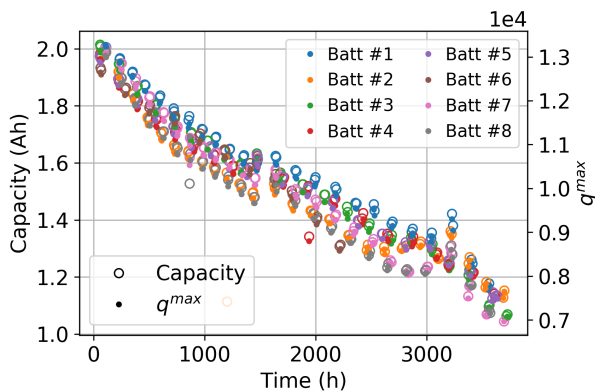
(a) Decrease in q^{max} over time.
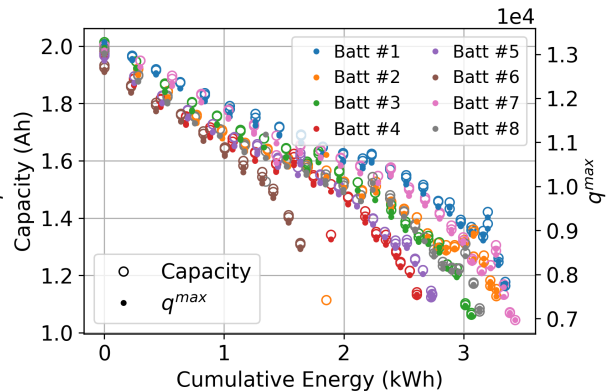
(b) Decrease in maximum capacity (left-vertical axis) and q^{max} (right-vertical axis) as a function of the cumulative energy.
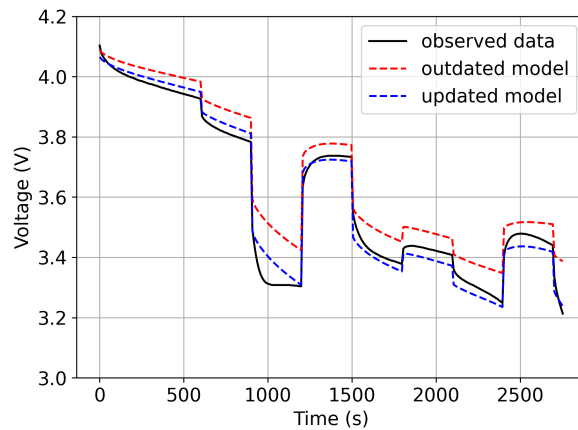
(c) Model voltage prediction under random input current loading before and after update the model parameters (q^{max}, R_0). Black curve is observed data, red curve is the prediction with outdated model parameters (RMSE 8.1e-02) and blue curve is the prediction with updated model parameters (RMSE 3.1e-02)

Figure 4.2: Variation of q^{max} as batteries accumulate cycles and its effects in the model voltage prediction.

In practice, incorporating the aging effect in the battery physics-informed neural network cell presented in Fig. 3.2b implies in incorporating the dependency of q^{max} and R_0 with respect to the cumulative energy drawn from the battery. Moreover, it is highly desirable to include uncertainty estimates in the forecasted values of the two parameters as a function of cumulative energy.

Table 4.1: Configuration of variational multi-layer perceptrons used for both q^{max} and R_0. ELU stands for exponential linear unit.

Layer	Mean # neurons	Activation	Variance # neurons	Activation
#1	4	elu	1	linear
#2	2	elu		
#3	1	linear		

In order to capture battery-to-battery variation due to aging, we use *variational*[1] multi-layer perceptrons [72–74] to build the models for q^{max} and R_0 as a function of cumulative energy. The variational models capture trends and dispersion of q^{max} and R_0 as a function of cumulative energy similarly to nonlinear regression models, in which parametric models describe the expected value, and discrepancy is modeled through Gaussian error with zero mean and unknown variance. The details for these models are given in Tab. 4.1. The Tensorflow Probability library is used for the ensemble model, allowing consistency with the model used to predict the voltage drop, and providing a unique solution capable of handling voltage drop predictions and aging at the same time. Figure 4.3 illustrates a comparison between the predictions from the variational models and the q^{max} used for training (similar performance was observed for R_0).

Aging and Voltage Drop Forecast with an Ensemble Model

Once q^{max} and R_0 starts being collected for a specific battery, one can build the proposed variational MLPs and use them to forecast q^{max} and R_0 even past the current cumulative energy for the battery. Nevertheless, the process is prone to large error if only few q^{max} and R_0 data points are

---

[1]With variational multi-layer perceptrons, each neural network parameter follows a Gaussian distribution, as opposed to being a deterministic value. In the optimization of parameters, the variational model uses a loss function that is the sum of negative log-likelihood and the Kullback-Leibler divergence [71], which estimates the difference between a reference distribution and the posterior distribution.
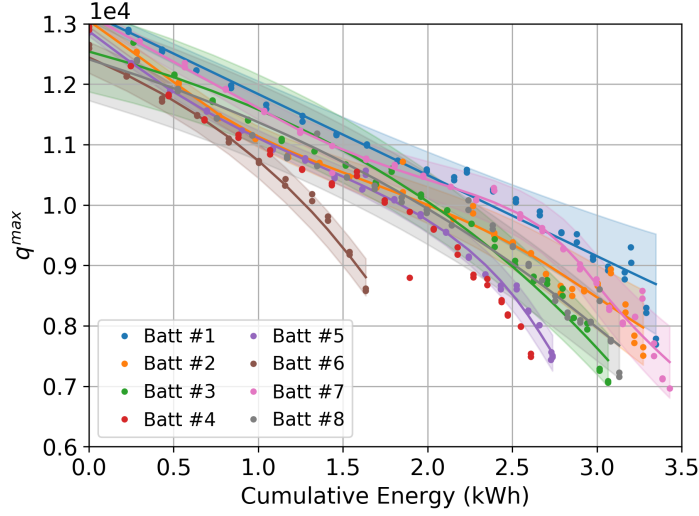
Figure 4.3: q^{max} and predictions from variational models as function of cumulative energy. Shaded areas denote the 95% prediction interval.

available.

Alternatively, predictions for a new battery can be improved if sub-population of batteries have already been used to collect a set of q^{max} and R_0. In fact, the set of variational MLPs built for each battery in the available sub-population can be seen as priors for the new battery being monitored. Bayesian update can be performed on the weights and biases of each of each model in the set. In addition, in order to benefit even further from the available set of models, we propose creating an ensemble model as a weighted average [68]. For example, for q^{max}:

$$\hat{q}^{max}(E) = \sum \omega_i q_i^{max}(E) \quad \text{and} \quad \omega_i = \frac{\mathbf{\Sigma}^{-1}\mathbf{1}}{\mathbf{1}^T\mathbf{\Sigma}\mathbf{1}}, \tag{4.1}$$

where $\Sigma_{ij} = \mathbf{e}_i^T \mathbf{e}_j$ and $\mathbf{e}_i$ is a vector with a measure of error. In this work, the prediction error between the i^{th} model in the ensemble and the observed data points is used for the battery that the ensemble is been built for.

39

The ensemble weights can be computed based on the predictions coming from the sub-population priors. Nevertheless, there are benefits in using ensemble weights out of the posterior models (particularly as the collected set of q^{max} and R_0 for the new batteries starts capturing its specific behavior). We argue that, over time, enough q^{max} and R_0 pairs will be collected and the effect of the sub-population information is reduced.

In order to validate the approach, an ensemble model with 7 out of the 8 batteries subject to random loading is used to predict the aging of the battery left out (battery #4 in this case). It is assumed that 6 data points in the (E, q^{max}) plane have been collected, which means that battery #4 has already been discharged 6 times. Then Eq. 4.1 is used to generate weights for the ensemble model. For illustration purposes, ensembles were built using both prior and posterior variational MLPs and show them hereafter. Figure 4.4 summarizes the weight-distributed and the forecasted q^{max} for battery #4 using this approach. Figure 4.4a shows in orange the q^{max} and respective prior variational MLP predictions for three batteries in the sub-population. It also shows the q^{max} used to obtain the posterior variational MLP predictions (in black and green, respectively). The future points of the battery being monitored (gray dots) are shown there for illustration only. Figure 4.4b shows the weights associated with each variational MLP comming from the sub-population of batteries. The weights indicate the similarity between each model and the current data points from battery #4 (left out for testing). Finally, Fig. 4.4c illustrates the forecasts before (orange lines) and after (green line and colored area) the Bayesian update. From left to right, the number of available points in the Bayesian update increases (from 6 to 16 to 26 points). As expected, reduction in forecast uncertainty is observed. This is true particularly when moving from 6 to 16 points. The ensemble model forecast varied less when moving from 16 to 26 points (indicating model convergence). The blue thick line and shaded area show predictions of the aging parameter q^{max} in a pure data-drive regression fashion, ignoring data from other batteries of the same fleet. The "observation only" approach suggests the variability of the results when ignoring prior information
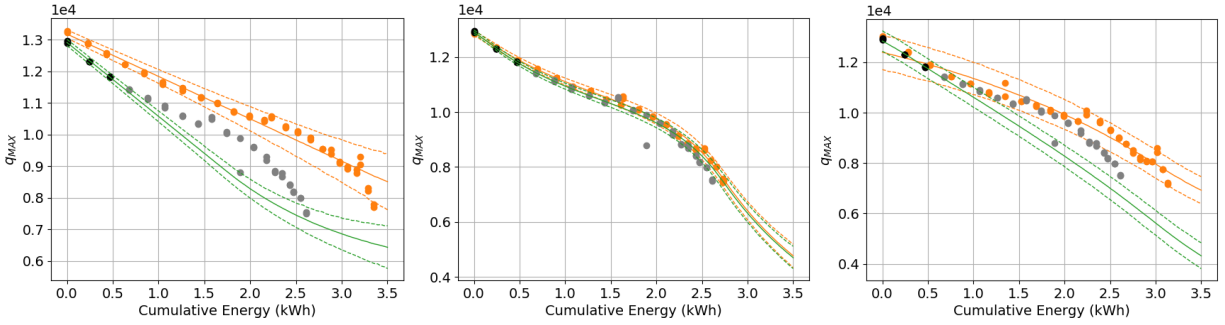
(a) q^{max} and individual prior and posterior variational MLP predictions. Information about batteries #1, #5, and #8 is shown in orange. Dashed lines denote the 95% prediction interval.
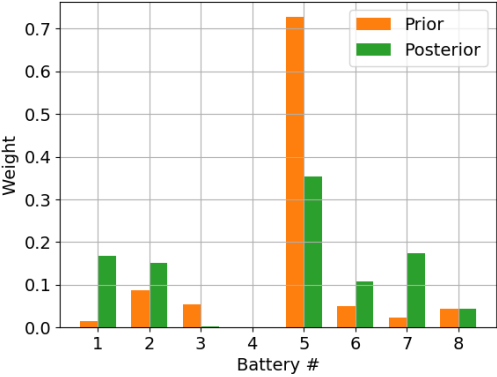
(b) Ensemble weights. Battery #4 is the one left out for validation.
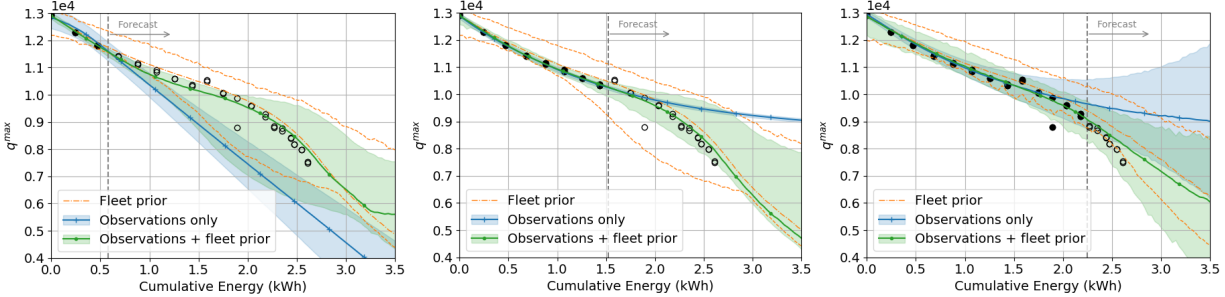
(c) Ensemble forecast for battery #4. From left to right, 6, 16, and 26 values of q^{max} were used in the Bayesian update and ensemble weight computation.

Figure 4.4: Example of q^{max} forecast using the ensemble of variational MLPs. Orange and green represent priors and posteriors, respectively. Black and gray dots are available and unknown (for the sake of the example) q^{max} values for the new battery, respectively. Shaded areas and dashed lines denote the 95% prediction interval.

from the fleet.

The ensemble model provides probabilistic estimates of the aging proxies at any given point in the future, and such values can be used to predict future voltage discharge cycles far-ahead in the battery life. This approach allows us to obtain predictions with uncertainty estimates of what the future discharge curve will look like after the battery reaches a certain aging level. As done previously, battery #4 is used for testing the method. After collecting data points of (E, q^{max}) and (E, R_0), the posterior ensemble model is used to forecast the values of q^{max}, R_0 and their uncertainty from the last collected cumulative energy all the way to $E = 2.5\,\mathrm{kWh}$. With that, the voltage discharge curve is predict assuming that the loading profile is known. Figure 4.5 illustrate some of the results. Figure 4.5a shows forecasted reference discharge cycles, which means that current held constant at 1 A. Figure 4.5b shows forecasted random discharge cycles (see top-right panel of Fig. 4.1 for current profile). All forecasts are made for $E = 2.5\,\mathrm{kWh}$. From the left to the right panel, data was collected up to when $E = 0.5\,\mathrm{kWh}$, $E = 1.5\,\mathrm{kWh}$, and $E = 2.25\,\mathrm{kWh}$. Results obtained with the variational model using observations from battery #4 only are compared against results obtained with the ensemble model using fleet priors. As expected, the predictions from the model that use only observations are considerably worst when compared to the ensemble model. In both cases (with or without fleet priors), the prediction of the voltage drop changed as more data points are collected from the aging of the battery; and therefore, voltage discharge forecasts improved. The expected value of the discharge curve predicted using fleet priors is relatively accurate even with only 6 data points were collected. The uncertainty on future values of q^{max} and R_0 propagates into large uncertainty bands on the future values of V, as expected (see in Fig. 4.4 how uncertainty in q^{max} varies with the number of data points).

(a) Reference discharge.

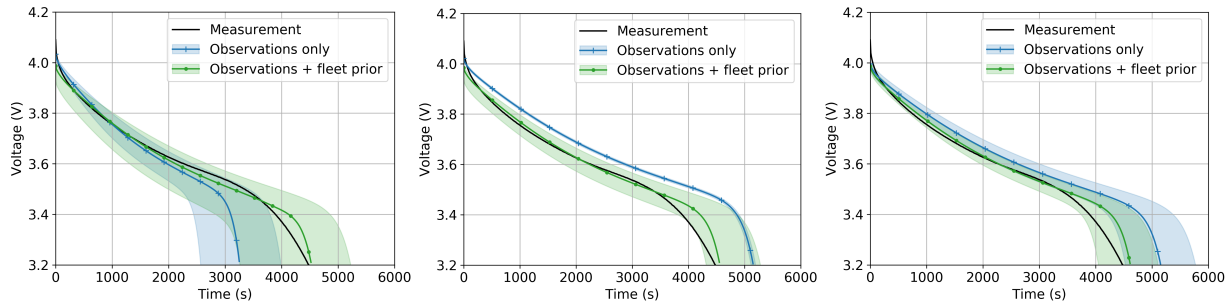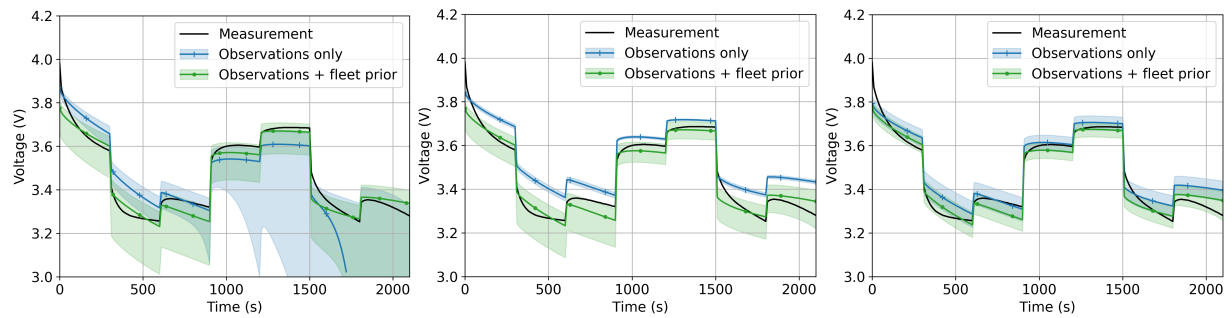(b) Random-loading discharge.

Figure 4.5: Forecast of voltage discharge curves for battery $\#4$. From left to right, 6, 16, and 26 values of q^{max} were used in the Bayesian update and ensemble weight computation (see Fig. 4.4). While all forecasted predictions are done at $E = 2.5$ kWh, models were built at $E = 0.5$ kWh (6 points), $E = 1.5$ kWh (16 points), and $E = 2.25$ kWh (26 points). Shaded areas denote the 95% prediction interval.

CHAPTER 5: DEALING WITH COMPLETE AND CENSORED DATA THROUGH SPECIALIZED LOSS FUNCTIONS AND VARIATIONAL MODELS

Based on our analysis shown in the previous chapters, the capability of the hybrid model to predict discharge curves at different aging conditions is heavily influenced by q^{max} and R_0. Since their behavior is strongly related with the residual capacity of the battery, we use these two parameters as proxies to represent battery aging. To achieve that goal, we utilized the constant discharge cycles performed at low current draw to estimate the residual capacity of the battery and improve predictions of future capacity degradation. However, constant discharge cycles are performed in laboratory environments and therefore are not ideal for operators, who should pull the battery out of commission to perform such tests.

This chapter addresses the need for such constant discharge cycles and shows how the degradation model based on fleet-wide data can be updated using censored data. This data is the estimated used capacity from random loading discharge cycles.

We propose a Bayesian approach to build a capacity degradation model that can handle data coming from both complete and incomplete discharge cycles. We fuse the reduced-order models based on first principles with modern deep neural networks and compensate for the reduced number of data points with prior information coming from a fleet of similar batteries. The proposed approach can be used for model updating and does not require constant discharge curves, which are the standard for estimating residual battery capacity. This is attractive to battery operators, as model updating can be performed without decommissioning the battery. In addition, this approach can handle battery-to-battery variation, which can happen due to factors such as inherent variability of manufacturing, initial internal damage, etc. We built an indicator of model prediction accuracy
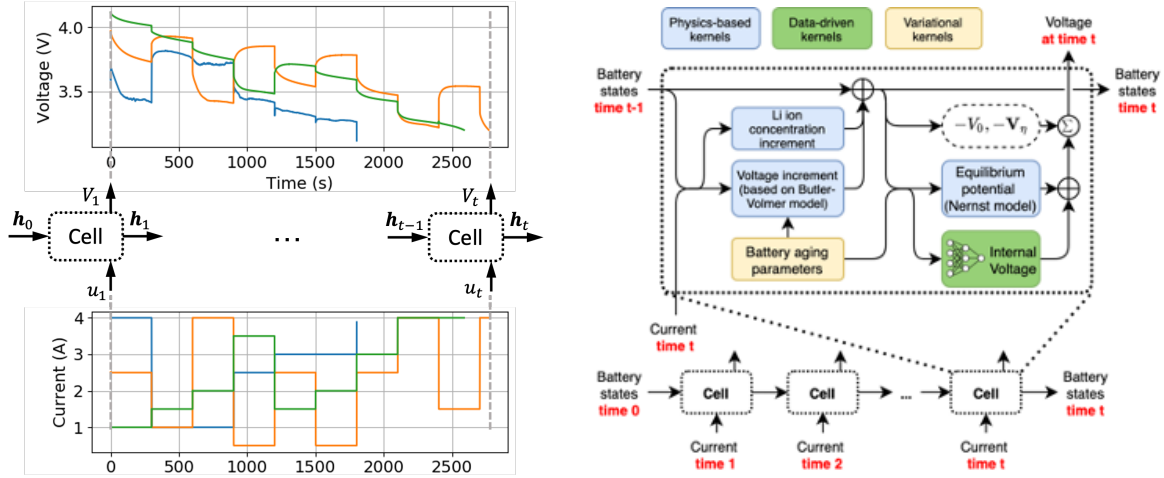
Figure 5.1: Battery hybrid RNN cell concept (left panel) and RNN implementation for time series prediction of output voltage given the input current (right panel).

using the Kullback-Leibler divergence of the energy drawn error distribution. By so doing, we can track the model's prediction ability so that another constant discharge cycle for the specific battery can be requested to improve the model's predictive performance.

Battery Aging as a Variational Multi-layer Perceptron

In this chapter, battery aging will be tracked by monitoring the decay of its capacity over the useful life. This is illustrated in Fig. 5.2a, which shows that capacity reduces as individual batteries accumulate used energy. We use a pair of *variational* MLPs (vMLPs) [72–74] to model battery capacity, C, as a function of cumulative energy, E:

$$C(E) \sim \mathcal{N}\left(\mu_C(E), \sigma_C^2(E)\right) , \; \mu_C(E) = \text{vMLP}_{\mu_C}\left(E; \mathbf{w}_{\mu_C}, \mathbf{b}_{\mu_C}\right) ,$$

$$\text{and } \sigma_C(E) = \text{vMLP}_{\sigma_C}\left(E; \mathbf{w}_{\sigma_C}, \mathbf{b}_{\sigma_C}\right) ,$$

(5.1)

where $\mathcal{N}(.)$ denotes the normal distribution defined by mean μ_C and variance σ_C^2; and $\mathbf{w}_{\mu_C}$, $\mathbf{b}_{\mu_C}$, $\mathbf{w}_{\sigma_C}$ and $\mathbf{b}_{\sigma_C}$ are the vMLP hyperparameters (in variational multi-layer perceptrons; hyperparameters follows a normal distributions; as opposed to being a deterministic values).

A major benefit of the formulation presented in Eq. (5.1) is that we can use both constant and randomized discharge data (Fig. 3.3) in the optimization of the vMLP hyperparameters. Constant
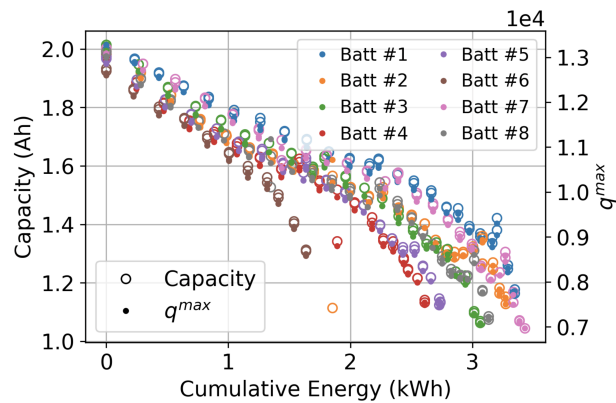
(a) Maximum capacity (left y-axis) and q^{max} (right y-axis) as a function of the cumulative energy.
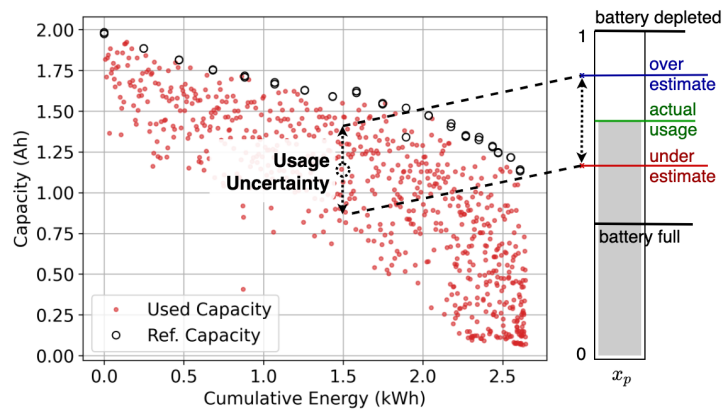
(b) Battery capacity used from random-loading discharge curves (small red dots) and constant discharge tests (large black dots)

Figure 5.2: Implicit relationship between capacity and q^{max} (a) and dispersion of used capacity from randomized cycles against residual capacity (b).

discharge data is typically performed in a controlled laboratory setting; and therefore, it is of extremely high quality. However, these experiments take the batteries out of service and can be expensive to obtain across an entire fleet. Randomized discharge is essentially regular usage data. While this data does not allow for direct assessment of battery capacity, it is readily available after every cycle. Figure 5.2b shows the capacity against cumulative energy for battery #4. The large black dots represent the constant discharge tests at constant loading; while the small red dots are the discharge curves under random loading conditions. It is important to notice that constant discharge provides complete data, as it allows assessment of the actual battery capacity. Conversely, random discharge provides censored data, as it only allows the assessment of the used capacity. In other words, with random discharge, although the actual capacity is not measured, it is known to be above the measured used capacity.

In order to accommodate both complete information (constant discharge) as well as censored data (random discharge), the sum of negative log-likelihood and the Kullback-Leibler (KL) divergence [71] is used as loss function in the optimization of the variational multi-layer perceptron hyperparameters:

$$\Lambda = \Lambda_{NLL} + \Lambda_{D_{KL}} \ . \tag{5.2}$$

The negative log-likelihood contribution, Λ_{NLL}, is computed using observed capacity (constant discharge) and used capacity (random discharge):

$$\Lambda_{NLL} = -\sum_i \log\left(\phi_{C,i}\right) - \sum_j \log\left(1 - \Phi_{C,j}\right) ,$$

$$\phi_{C,i} = \phi_C\left(c_i; \mathbf{w}_C, \mathbf{b}_C, \sigma_C\right) , \qquad \text{and} \qquad \Phi_{C,j} = \Phi_C\left(c_j; \mathbf{w}_C, \mathbf{b}_C, \sigma_C\right) , \tag{5.3}$$

where $\phi_C(.)$ and $\Phi_C(.)$ are the probability density and cumulative distribution functions of the Normal distribution, as assumed in Eq. (5.1); and c_i and c_j are the complete and censored data, respectively.

The KL-divergence contribution, $\Lambda_{D_{KL}}$ is computed across the vMLP layers using the following definition:

$$\Lambda_{D_{KL}} = \sum_i \sum_k \log(\phi_{A,k}) - \log(\phi_{B,k}) \,,$$

$$\phi_{A,k} = \phi_A\left(c_i, \mathbf{w}_{C,k}, \mathbf{b}_{C,k}; \mu_{\mathbf{w}_C,\mathbf{b}_C}, \sigma_{\mathbf{w}_C,\mathbf{b}_C}\right) \,,$$

$$\text{and} \quad \phi_{B,k} = \phi_B\left(\mathbf{w}_{C,k}, \mathbf{b}_{C,k}; \mu_{\mathbf{w}_{C,0},\mathbf{b}_{C,0}}, \sigma_{\mathbf{w}_{C,0},\mathbf{b}_{C,0}}\right) \tag{5.4}$$

where $\phi_A(.)$ is the likelihood of the vMLP weights and biases; and $\phi_B(.)$ is the probability density function of the prior distribution of the vMLP weights and biases.

Battery-to-battery variation due to aging will be tracked through parameters q^{max} and R_0, as detailed in the hybrid physics-informed neural network of Fig. 5.1. Figure 5.2a shows the high degree of agreement between the q^{max} parameters with the actual battery capacity. R_0 is also strongly correlated to C but it has not been shown for the sake of brevity. Therefore, models for q^{max} and R_0 are easily built using the capacity model through:

$$q^{max}(E) = \alpha\, C(E) \,, \quad \text{and} \quad R_0(E) = \frac{C(E)}{\gamma(E)} \,, \tag{5.5}$$

where α is a simple scaling factor; and $\gamma(E)$ is a linear model. The initial capacity is estimated on a battery-to-battery basis using the first constant discharge cycles in the data set. Then, R_0 and q^{max} are initialized using Eq. (5.5). This approach assumes that each new battery undergoes at least one constant discharge cycle before entering the field.

Priors from Fleet Information and Strategy for Continuous Update

We propose using available fleet-wide information to build priors for the capacity model described in Eq. (5.1). This approach has two major advantages. Firstly, it allows reliable forecast of capacity for unseen mission profiles that new batteries will be subjected to in the future; even if these batteries were not used to build the fleet-wide priors. The ability to forecast capacity, and therefore, q^{max} and R_0, allows the use of our hybrid variational physics-informed neural network to predict and forecast battery discharge as well. Secondly, fleet priors can facilitate estimating model parameters for new batteries when constant discharge cycles are not abundantly available.

We propose the following steps to build fleet wide priors (FWP):

FWP.1: Given a set of batteries representative of a fleet, collect historical actual capacity data by performing a series of constant discharge cycles (random discharge cycles in between constant discharge cycles should be used to accelerate battery aging). This step generates the data illustrated in Fig. 5.2a.

FWP.2: For each battery in the set (step *FWP*.1), build capacity models using the formulation described in Eq. (5.1). The collection of models capture the capacity decay over usage accounting for battery-to-battery variation.

FWP.3: Use the collection of models as prior information for new (unseen) batteries.

For a new battery, not included in the set used to generate fleet priors, we assume that a initial constant discharge is performed prior to deploying the battery into service. Once the battery is in use, different duty cycles will generate a series of random discharges. With that, used capacity data on a cycle-by-cycle basis is plentiful. We propose the following steps to build and continuously update the battery-specific model:

BSM.1: Use fleet-wide information as prior (*FWP*.1 to 3).

BSM.2: As new data becomes available (either random or constant discharge cycles), use Eq. (5.2) to continuously update the battery-specific aging model described in Eq. (5.1).

BSM.3: When new data is dominantly formed by random discharge cycles, we propose balancing the contribution from individual batteries in the fleet with an ensemble model as a weighted average [68]:

$$C_{\text{WA}}(E) = \sum \omega_i C_i(E) \quad \text{and} \quad \omega_i = \frac{\mathbf{K}^{-1}\mathbf{1}}{\mathbf{1}^T\mathbf{K}\mathbf{1}}, \tag{5.6}$$

where $K_{ij} = \Lambda_{NLL,i}^T \Lambda_{NLL,j}$ and $\mathbf{\Lambda_{NLL}}$ is a vector with log-likelihood values - see Eq. (5.3).

Along the same lines of what we propose to building battery-specific models for aging, we can use the resulting probabilistic hybrid model to guide acquisition of new constant discharge data. Again, suppose a new battery is first cycled through a constant load cycle; and then, put into service where it is submitted to a large number of random discharge cycles. We propose the following steps to detect the need for new constant discharge data:

NCD.1: Build and constantly update the battery-specific aging model using the previously described steps *BSM*.1-3.

NCD.2: After every random discharge compute the distribution of energy for the given duty cycle. The estimated energy error distribution is obtained by subtracting the actual value from the estimated ones.

NCD.3: Track the divergence of the energy error distribution using metrics such as the KL-divergence and others.

NCD.4: Collect new constant discharge data as the divergence metrics cross tolerable thresholds.

NCD.5: Update the battery-specific model using steps described in *BSM*.1-3.

Effectively, these steps leverage the battery-specific aging model to detect when the hybrid physics-informed neural network predictions are drifting ever so slightly. This drift indicates the need for new constant discharge cycles, which will then be used to perform a major update in the model. This happens due to the fact that complete data has a much stronger contribution to the loss function, Eq. (5.2), when compared to censored data.

<div align="center">Results and Discussions</div>

Our implementation is all done in TensorFlow[1] (version 2.3) using the Python application programming interface. The uncertainty quantification and probabilistic modeling is performed using the TensorFlow Probability library [2] (version 0.11.0). The interested reader can find an in-depth discussion of how to implement hybrid recurrent neural networks in TensorFlow in [75]. Our hybrid model was validated against the physics-based model, which is publicly available in the NASA's *Prognostics Model Library* [64]. The next sections will detail how the model is fitted and validated as well as its use for random discharge prediction and modeling of battery aging.

<div align="center">Fitting and Validating the Hybrid Variational Physics-Informed Neural Network</div>

Table 5.1 details the design of the two MLPs and table 5.2 the variational MLP used in the hybrid model. With regards to the non-ideal voltage MLPs, similarly to the pure physics model, a linear model is used for the negative side, which is a MLP with one layer, one hidden unit and no activation function. The non-ideal voltage curve for the positive side is more complex, and

---

[1]`www.tensorflow.org`

[2]`www.tensorflow.org/probability`

Table 5.1: Configuration of the multi-layer perceptrons used to approximate the behavior of the non-ideal voltage. The activation functions used are: hyperbolic tangent (tanh) and linear (lin).

Layer	$V_{ni,p}$ # neurons	activation	$V_{ni,n}$ # neurons	activation
#1	8	tanh	1	linear
#2	4	tanh		
#3	1	tanh		

Table 5.2: Configuration of the variational multi-layer perceptron used for $C(E)$ – see Eq. (5.1). The activation functions used are: exponential-linear unit (elu) and linear (lin).

Layer	$\mu_C(E)$ # neurons	activation	$\sigma_C(E)$ # neurons	activation
#1	4	elu	1	linear
#2	2	elu		
#3	1	linear		

therefore a MLP with 3 layers (8 hidden units in the first layer and 4 in the second hidden layers) with hyperbolic tangents as activation functions is used. All hidden units and output units have biases. Both MLPs take the mole fractions x_n, x_p, as inputs, and return the corresponding non-ideal voltages. The larger MLP for the estimate of $V_{ni,p}$ has 57 trainable parameters, while the MLP for the negative-side non-ideal voltage has only two trainable parameters (slope and bias of the linear model). The variational MLP used to model the location parameter of the capacity C as a function of E has three layers and 42 trainable parameters, while the scale parameter has only one layer, and as such, 4 trainable parameters. All four models were trained with the Adam optimizer [66] set with a learning rate of 5×10^{-3} and 3000 epochs. All other Adam parameters were kept equal to their default values.
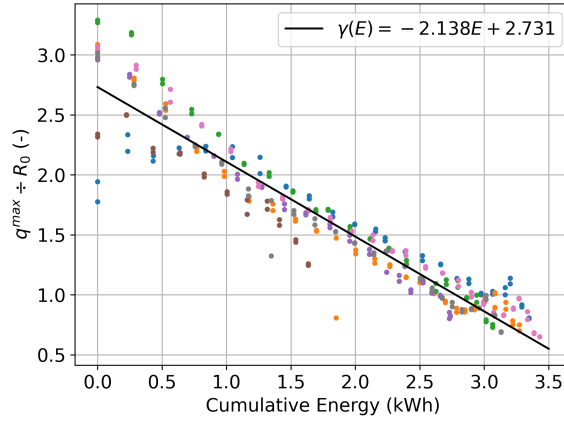
We observed that the relationship between the non-ideal voltage $V_{ni,i}$, $i = \{n, p\}$ and the concen-
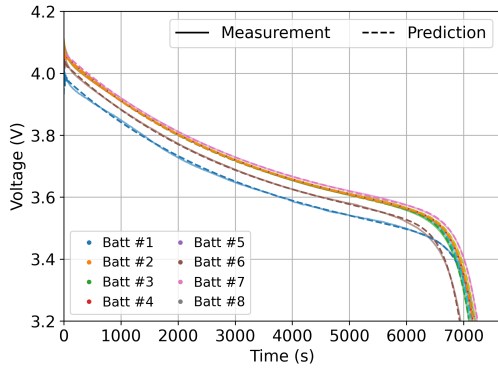
(a) Multi-layer perceptron model for internal voltage – predictions versus test data
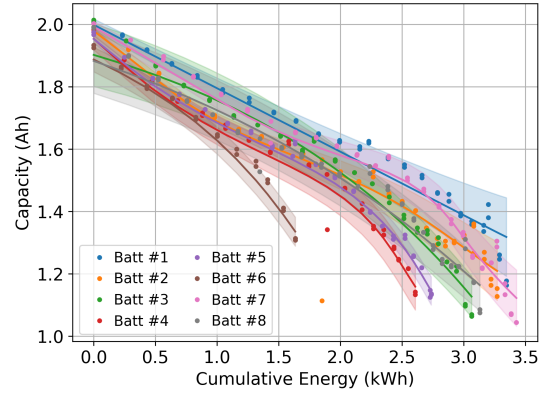
(b) Variational multi-layer perceptrons for battery capacity – results for all 8 batteries in the data-set.

(c) Model that helps mapping capacity into R_0, see Eq. (5.5).

Figure 5.3: Modeling elements of the hybrid variational physics-informed neural network.

tration overpotential $x_{s,i}, \; i = \{n, p\}$ is either not affected by aging or the effect appears negligible for the purpose of prognosis. Therefore, we utilized the first constant-loading discharge curves to train the non-ideal voltage models as well as estimate baseline values for q^{max} and R_0. Figure 5.3a shows how the predictions of the models from the trained MLPs compare against measured voltage discharge curves. Next, we built battery-specific vMLP models that correlate C to E, as illustrated in Fig. 5.3b. The plot highlights the importance of capturing battery-to-battery aging
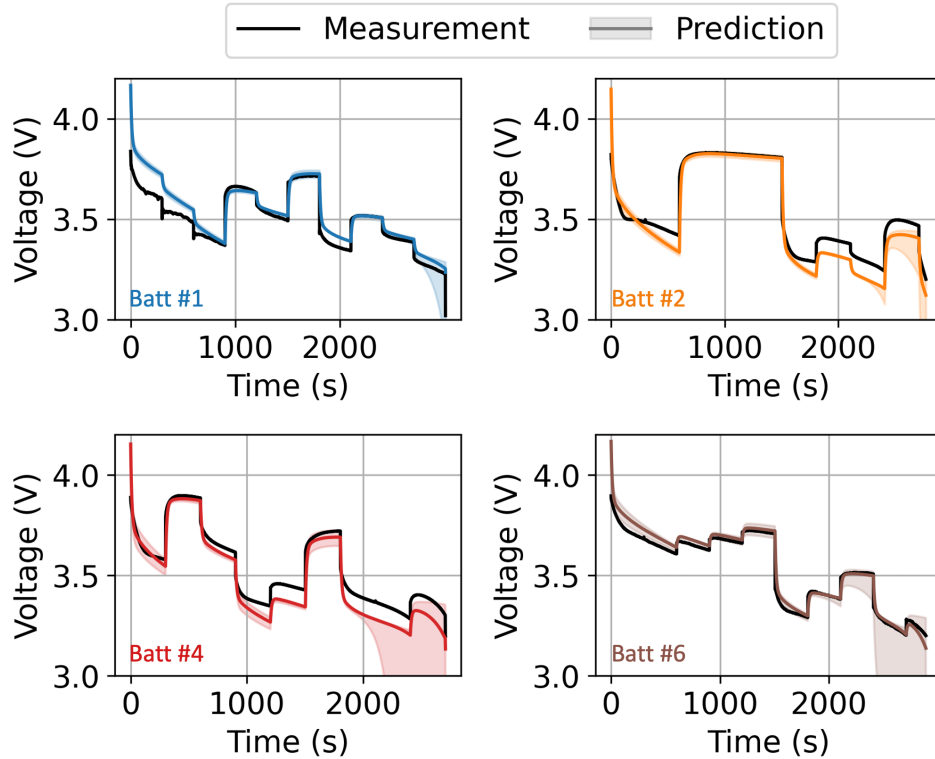
53

Figure 5.4: Random walk discharge predictions against measured data.

variability (later in this chapter, we will use the ensemble to predict the fleet behavior using future values for E). Once the capacity models are built, we used a scaling factor to estimate q^{max} from C, and a linear model to map q^{max} to R_0, as shown in Fig. 5.3c. With the pieces of the model properly fitted, we can use the hybrid variational physics-informed neural network to predict random discharge cycles. Figure 5.4 shows the predictions for different batteries in the set and highlights the uncertainty bounds in prediction, a key advantage of the proposed approach. We intentionally showed predictions for batteries across the capacity distribution from the fleet.

Updating Battery-specific Aging Model with Random Discharge Data

Now, we will focus on the case in which a new battery is inserted in the fleet. As opposed to relying only on constant discharge cycles to build the model for capacity as a function of cumulative energy; we will consider that, at first, only one constant discharge cycle is available. Once the battery is put into service, the only new information available are the random discharge cycles. As illustrated in Fig. 5.5a, this imposes a challenge for the use of the hybrid models for long-term forecast. As the battery ages, predictions for discharge cycles start drifting away from actual values. This is a clear indication that q^{max} and R_0 need to be updated.

As discussed in the previous sections, we compensated for the lack of regular constant discharge data using a Bayesian model update approach. This part of the study will be performed using a cross-validation approach. We extract one battery from the fleet and use the remainder to build priors from fleet-wide information in the form of a set of vMLPs for capacity vs. cumulative energy, as shown in Fig. 5.3b. Then, we deploy the methodology discussed in *BSM*.1-3, and update the fleet priors with battery-specific random discharge data. The update of the capacity model for battery #4 is illustrated in Fig. 5.5b. The prior is a simple ensemble of all batteries in Fig. 5.3b, except battery #4. The available data at 1 KWh is composed of one constant discharge at 0 KWh and all random discharge cycles run up to 1 KWh. The updated model is obtained using Eqs. (5.2) through (5.6) and can be used to forecast capacity as a function of cumulative energy (posterior in Fig. 5.3b). The updated model can be used to track current cycles, and more importantly, forecast future cycles. Figure 5.5c shows the improvement in discharge forecast at 2 KWh from the models before and after the update.
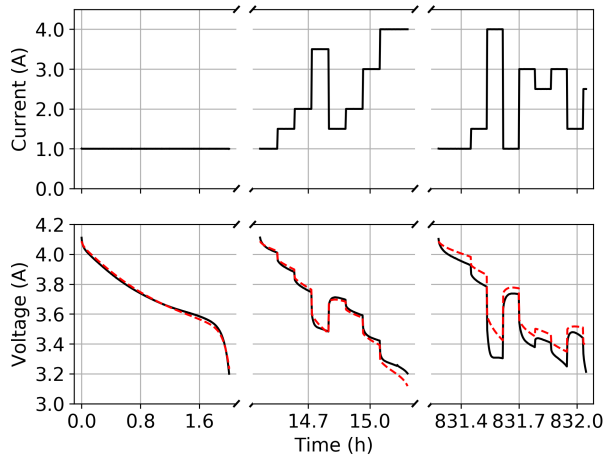
(a) Battery discharge predictions vs. measured data over time for different duty cycles.
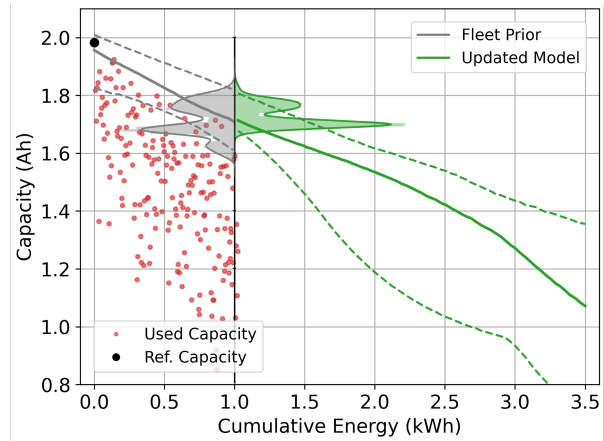
(b) Bayesian update of aging model with one constant discharge and several random discharge data points.
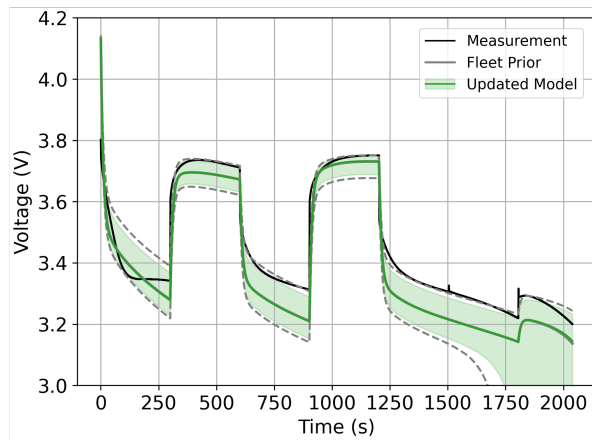
(c) Forecast at 2 KWh before and after model update.

Figure 5.5: Model drifting due to aging and its update using fleet-wide prior, an initial constant discharge cycle, and several random discharge cycles.

Adding New Constant Discharge Data

While updating the battery-specific capacity models with random discharge data is convenient in terms of data acquisition; it is also limited in terms of how much the model can improve in case

predictions are not satisfactory for a a specific battery. Therefore, it is expected that over time, new constant discharge data will be necessary to re-adjust the models.

We proceed the investigation using a cross-validation approach. The data from the battery that was extracted is used to study the use of KL-divergence of the energy error as an indicator for model predictivity allowing us to (a) confirm that we can use the fleet-wide information as prior distribution for the Bayesian update; and (b) evaluate the KL-divergence of energy error as an indicator of the health of the model.

We used the steps presented in *NCD*.1-5 to check the ability of the model to track battery-specific aging (i.e., the "health of the model") and then suggest when it is time to add a new constant discharge cycle. Figure 5.6a shows the estimated energy error distribution obtained at one given random discharge cycle of battery #4 when its cumulative energy was already at 0.5 KWh. Each realization in the bottom histogram is obtained by computing the area under the curve of power vs. time (power is obtained by multiplying voltage by current, in which case the our hybrid model generates a series of different realizations). In other words, the histogram represents the uncertainty in used energy coming from the hybrid model predictions. The error in energy prediction is obtained by subtracting the actual values from the predicted ones. This distribution can be tracked over time, as illustrated in the top panel of Fig. 5.6b. As long as the distribution remains stationary, the hybrid model is expected to be predictive for the particular battery. We use the KL-divergence between the present and initial energy error distributions as health indicator for the hybrid model. Once the KL-divergence passes above a threshold, we claim that the predictions given by the hybrid model are no longer representative of the actual system, as depicted in the bottom panel of Fig. 5.6b. Therefore, a new constant discharge cycle is needed to re-adjust the model due to the aging of the battery (updating the vMLP parameters).

Figure 5.7 shows the results for the capacity model as well as an example of forecast random
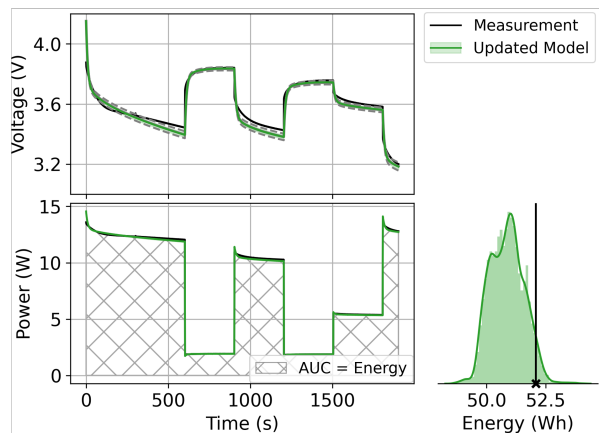
(a) Estimating the energy error distribution at 0.5 KWh.
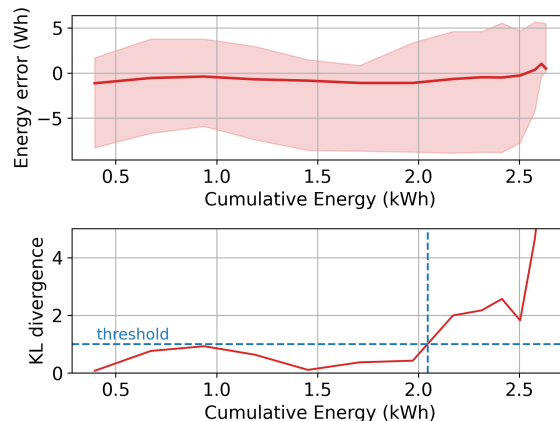
(b) Tracking energy error distribution and its KL–divergence vs cumulative energy.

Figure 5.6: Indicator of predictivity of the hybrid variational physics-informed neural network, using the energy error distribution – results for battery #4.

discharge cycle obtained after the model is updated with an additional constant discharge cycle. Figure 5.7a illustrates the capacity data from constant and random discharge cycles, the fleet-wide prior, and the posterior model resulting from the update using the process discussed in *BSM*.1-3. With the updated variational MLP plugged back into the battery-specific model, we can again accurately forecast future random discharge cycles, as shown in Fig. 5.7b.

The same process used to monitor the predictive capabilities of the battery-specific model can be repeated for all batteries in the set and results are shown in Fig. 5.8a. Overall, we observe that the KL-divergence appears to modulate until its value crosses a threshold around 1.0, before it starts to exponentially increase. Therefore, based on this cross-validation study, it is fair to suggest the a KL-divergence of the energy error of 1.0 for the battery family used in this research. Finally, we highlight that, as an indicator of how predictive the battery-specific aging model is, the KL-divergence will reduce after new constant discharge cycles are incorporated in the data set. Figure 5.8b illustrates how the metric varies over the battery #4 usage life. The trajectory with only
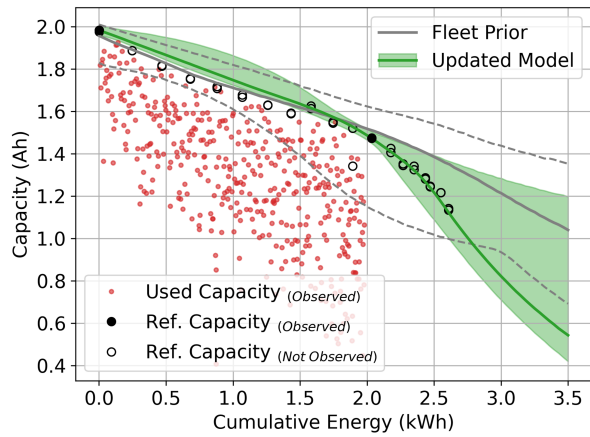
(a) Bayesian update of aging model with two con-
stant discharge and several random discharge data
points.
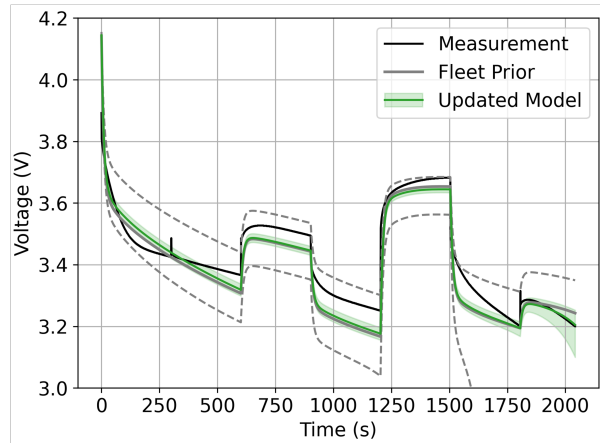
(b) Forecast at 2.5 KWh before and after model up-
date.

Figure 5.7: Model update using fleet-wide prior, an initial constant discharge cycle, several random discharge cycles, and one additional constant discharge cycle (added at 2 KWh as supported by Fig. 5.6).
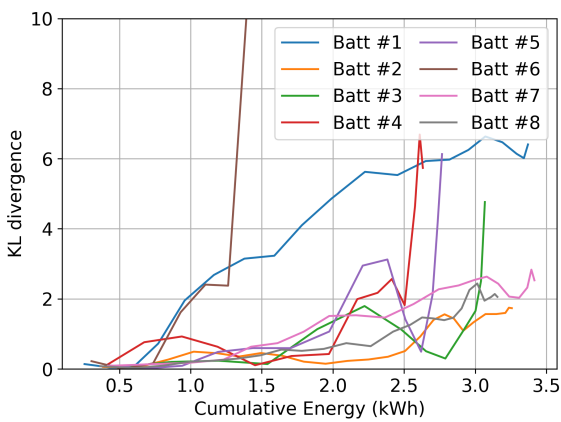
(a) Cross-validation study using 8 batteries.
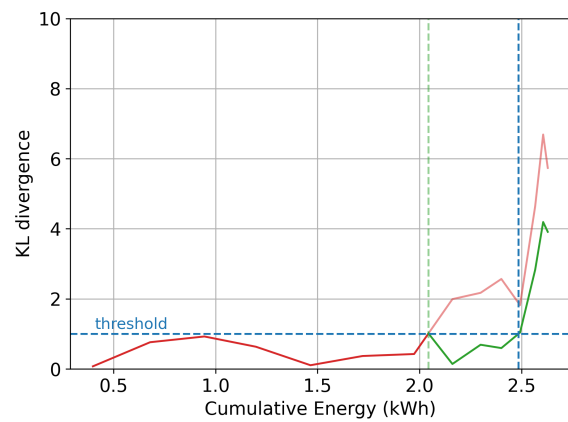
(b) Trajectory correction after a new constant dis-
charge is added (results for battery #4).

Figure 5.8: Study of KL-divergence metric across multiple batteries and after model update.

the first constant cycle and the several random discharge cycles is shown in in red. Without adding a new constant discharge cycle, the KL-divergence metric keeps growing. On the other hand, the green line shows the KL-metric when the model updated with data from a new constant discharge cycle added around 2.0 KWh. The KL-divergence metrics is immediately brought down to values well bellow 1.0 and, when it starts growing again, it becomes an indicator that yet another constant discharge cycle is needed. For battery #4, the suggestion of this third constant discharge cycle will happen around 2.5 kWh. At that point, the battery operation should consider whether it pays off to keep operating the battery given that capacity itself will be reduced from 2 Ah to to roughly 1.2 Ah (or 60% of the initial value).

CHAPTER 6: QUANTIFYING UNCERTAINTY IN AGING DUE TO UNKNOWN BATTERY USAGE THROUGH BAYESIAN UPDATE

The main motivation of this chapter is to offer estimates of remaining useful life even when the past battery usage is unknown. This scenario is likely to occur when prognosis and health management models are used to monitor legacy fleets of battery-powered vehicles. Operators, insurers, and service providers might have to serve vehicles with unknown previous operations, route structure, driver/pilot behavior, etc. Nevertheless, to control operations and maintenance costs, it is important to quantify battery degradation and estimate capacity fade over future missions. These models can also aid operators in optimizing fleet operation and allocation of assets across different mission mixes.

In this chapter, we assume that pre-existing data is available for at least a small fleet of the same battery model. While we do not require the available data to be large, it is important that the data is representative of the population, capturing battery-to-battery variation (as discussed in previous chapters). For the batteries lacking history of usage, we assume that we can record data for new and future missions, as well as specify when to acquire constant discharge data.

We then propose using Bayesian statistics to build models for the remaining useful life of these legacy batteries. As prior, we will use the variational hybrid physics-informed neural networks built on top of fleet-wide data. Then, we use sequentially collected constant discharge data to update the uncertainties about capacity versus cumulative energy for the specific battery. With the updated models, we can forecast battery performance degradation (fade in capacity) due to usage.

Bayesian Model Update for Battery Capacity as Function of Variation in Cumulative Energy

Figure 6.1a shows the history of capacity versus cumulative energy for the battery dataset used in this research. One could argue that constant discharge cycles, performed regularly since the battery is first deployed all the way to the end of its useful life, is a valid approach. However, in real-life applications, this might not necessarily be the case. For many legacy or already fielded batteries, the history of use might not be known and constant discharge cycles might not have been performed. Figure 6.1b illustrates a case in which the battery is known to have already accumulated usage, but how much is unknown. Therefore, even when plotting both reference (constant) and random discharge cycles, there is no practical way to stipulate what is the battery initial capacity. Nevertheless, in this research, we assume we can track battery capacity (and for the sake of our models, quantities such as q_{max} and R_0) versus the cumulative energy differential with respect to an arbitrary point:

$$C = f(\Delta e)\,; \tag{6.1}$$

where C is the battery capacity, Δe is the increment in cumulative energy from the moment the battery starts being monitored; and $f(.)$ is the unknown function that maps the cumulative energy increment into battery capacity (this is the function we want to build models for).

Figure 6.2 illustrates the fleet-wide probabilistic model for capacity as a function of cumulative energy (see Chapters 3 and 4) and show how a cumulative energy distribution can be obtained for a given observed capacity level. We argue that the large uncertainties in the cumulative energy distribution prevents it from being used for prognosis. However, we should be able to improve the uncertainty estimates as more data points, and particularly constant discharge cycles, are collected. In fact, here we propose using Bayesian statistics in order to update the uncertainty about battery capacity as a function of battery usage.

(a) Decrease in maximum capacity (left-vertical axis) and q^{max} (right-vertical axis) as a function of the cumulative energy.

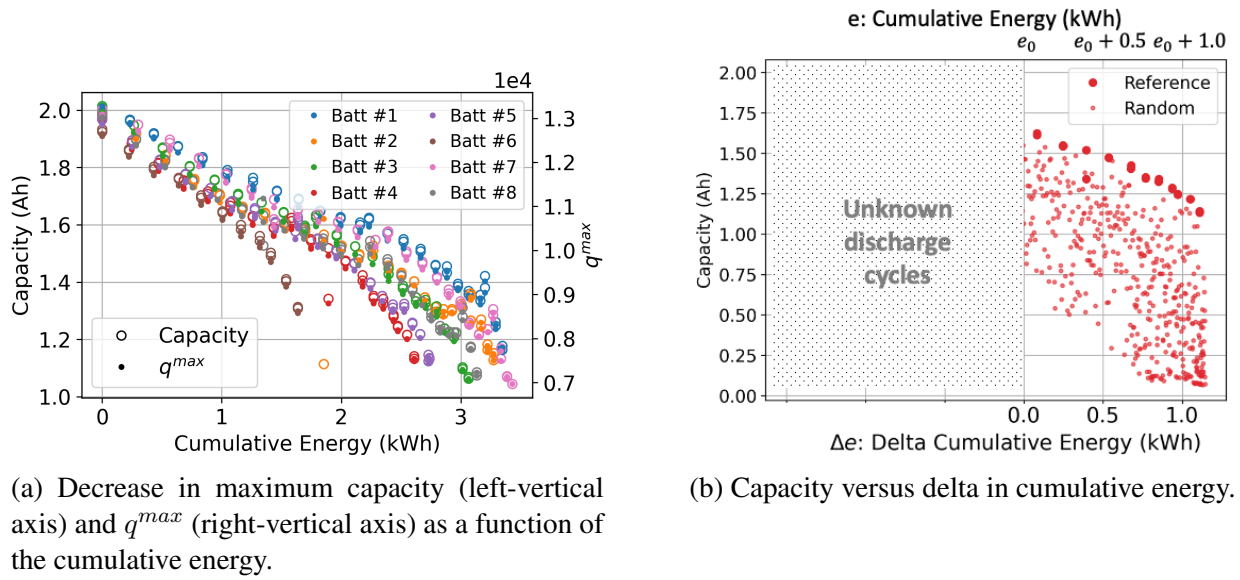(b) Capacity versus delta in cumulative energy.

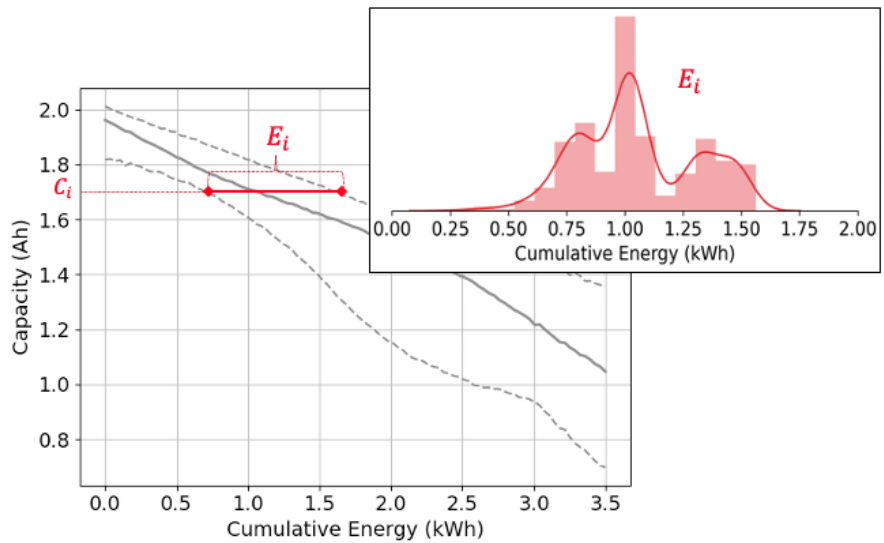Figure 6.1: Known versus unknown capacity versus battery usage history.

Figure 6.2: Fleet-based cumulative energy distribution at observed capacity level.

We start by defining the case in which an operator performs two reference discharge tests. The first one is performed when the battery starts being monitored, at unknown cumulative energy history. The second one is performed after the battery has gone through an arbitrary number of random discharge cycles. Figure 6.3a shows two observed capacity levels, c_0^* and c_1^*, coming out of constant discharge cycles, and the fleet-wide probabilistic model for capacity as a function of cumulative energy. With the prior uncertainty in cumulative energy modeled by the fleet-wide model, we can arrive at E_0 and E_1; the cumulative energy distributions at the c_0^* and c_1^*, respectively. However, since the battery has now been monitored; the increment in cumulative energy between the two constant discharge cycles is also known. Figure 6.3b illustrates the case in which that increment is $\Delta e = 0.189$ KWh. With this information; we can filter the uncertainty distribution in the fleet-wide capacity model such that $\Delta e = E_1 - E_2 = 0.189$. The result is shown in Fig. 6.3c.

The simple filtering of the fleet-wide model such that the constraint in the cumulative energy increment is respected ($\Delta e = 0.189$ kWh in this case) is already helpful in terms of building the battery-specific prognosis model. However, we propose using improve such models by using Bayesian statistics such that

$$p(\Delta E(C; \boldsymbol{\theta})) = \frac{L(\mathbf{D}|\boldsymbol{\theta}) \, p_0(\Delta E)}{\int L(\mathbf{D}|\boldsymbol{\theta}) \, p_0(\Delta E) d\boldsymbol{\theta}} \; ; \qquad (6.2)$$

where $\Delta E(.)$ is the random variable that defines the increment in cumulative energy as a function of C and $\boldsymbol{\theta}$; C is the battery capacity; $\boldsymbol{\theta}$ is the vector of parameters for the ΔE model; $\mathbf{D}$ is the set of observed data; which in this case is the set of observed capacity and cumulative energy increment pairs observed since the battery started being monitored; $p(.)$ and $p_0(.)$ are the posterior and prior distributions for $\Delta E(.)$; and $L(.)$ is the likelihood model.
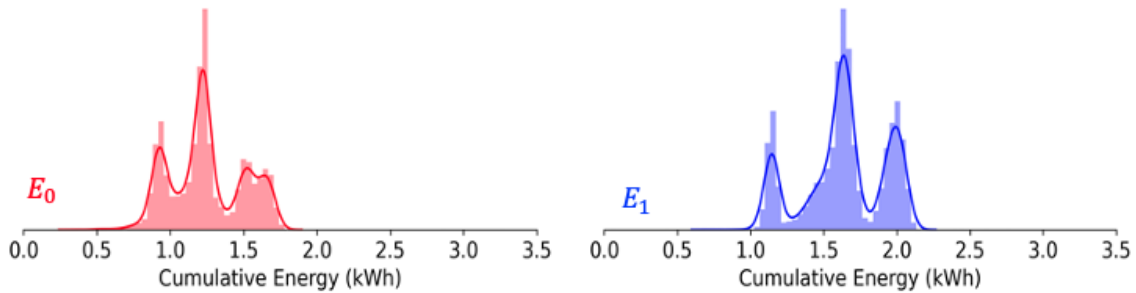
Equation (6.2) is known to be hard to solve due to the integral in the denominator. In order to sample the posterior distribution, here, we use numerical integration through particle filtering [63,
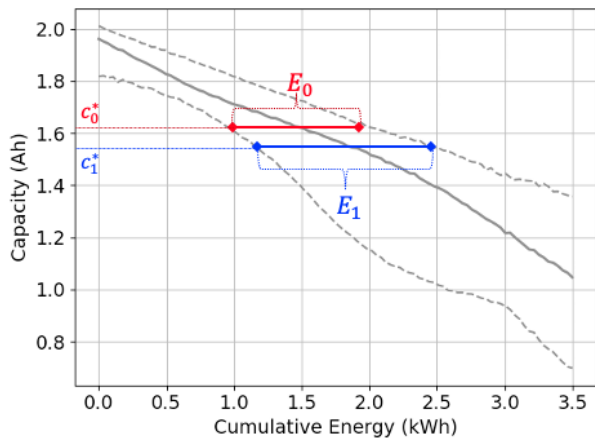
(a) Fleet-based capacity models and two observed capacity levels.
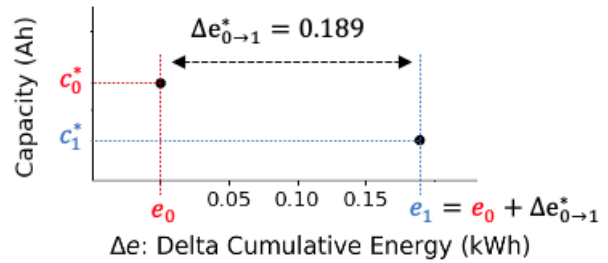
(b) Known offset between cumulative energy at observed capacity levels.

(c) Cumulative energy distributions at observed capacity respecting offset.

Figure 6.3: Manipulating fleet-wide prior capacity model with observed capacity at two levels.

76, 77]. For convenience, here, we assume the likelihood to be Gaussian. Our Bayesian update strategy is illustrated in Fig 6.4.

Results and Discussion

Figure 6.5 illustrates how the proposed approach perform when only two observations of constant discharge are available. Each subplot has two sets of images. The prior, likelihood, and posteriors for Δe are shown at the top, while the data and models for capacity versus Δe are shown at the
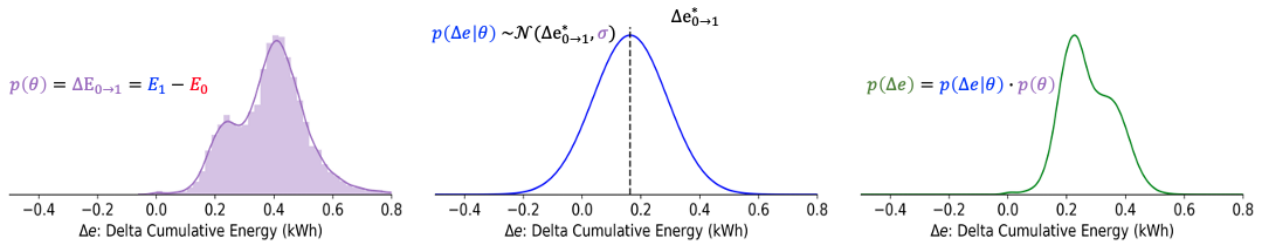
Figure 6.4: Applying Bayesian update to Δe.

bottom. Figures 6.5a and 6.5b show cases in which data is acquired early on in the useful lives of two different batteries, and Fig. 6.5c illustrates the case in which the data is collected half-way through the battery useful life. It is important to notice that, in all cases, we illustrate results out of cross validation. That is, the batteries shown in the Bayesian update were not in the set used to obtain the priors.

We can expect that when only two constant discharge data points are available the prior information is extremely important. Therefore, the posterior distribution, as seen in the posterior of the model for capacity versus Δe, is prone to perform best when data is observed around the prior distribution. In other words, since the prior captures the uncertainty in the overall population, the closer to the average behavior the data falls, the better the posterior model will be. This is not a limitation of our method; instead, it is a consequence of the Bayesian approach under very limited data. This is illustrated by the contrast between Figs. 6.5a and 6.5b.

A major main benefit of our proposed approach is that priors compensate for the lack of data (as previously discussed). However, when data becomes available, the contribution coming from the likelihood term substantially increases. Over time, when data is plentiful, the mild priors elucidated from the small population (in the cases illustrated here, only 7 batteries). Figure 6.6 illustrates cases in which data is sequentially added. Similarly to the Fig. 6.5, prior, likelihood, and posterior are shown at the top; and capacity model and data versus Δe are shown at the bottom.
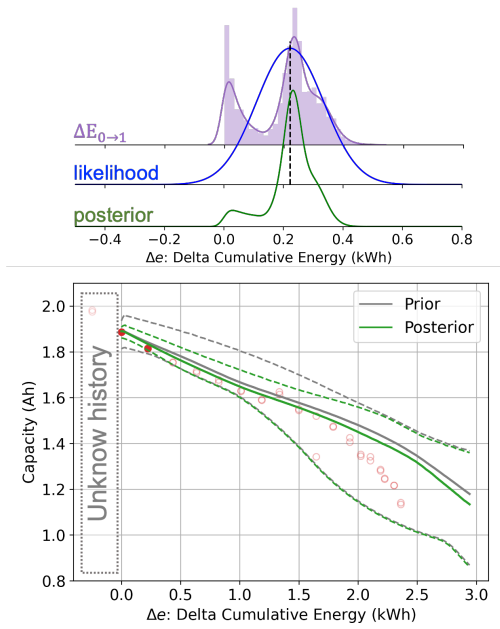
(a) Battery 4, early life.
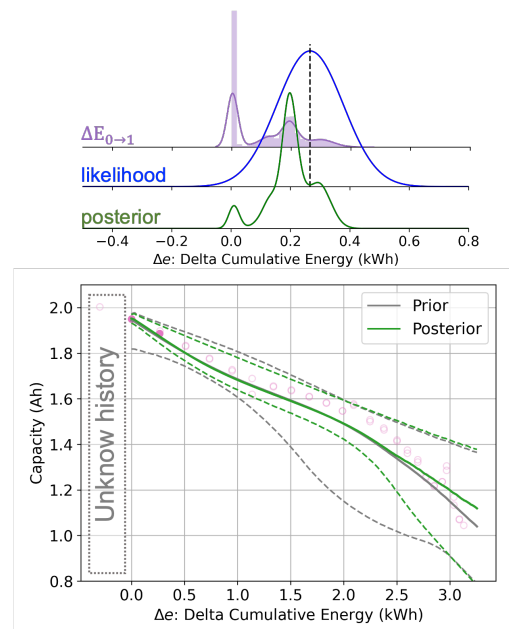
(b) Battery 7, early life.
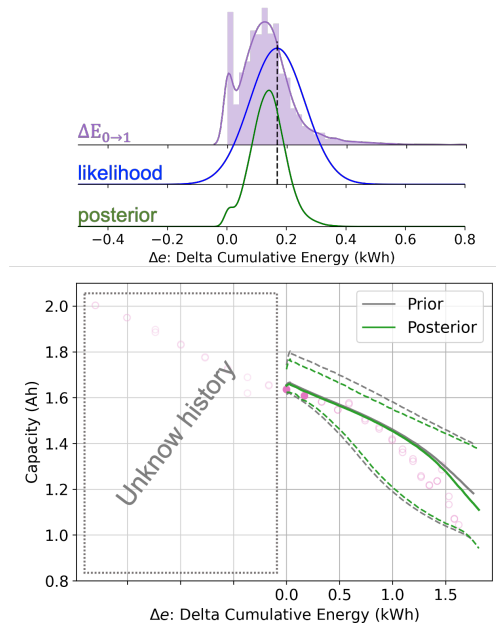
(c) Battery 7, mid life.

Figure 6.5: Battery-specific capacity vs Δe models updated with two observations.

While Fig. 6.6a shows a case in which data is sequentially added early on in the useful life of a battery; Fig. 6.6b shows a case in which data comes only half-way through the battery useful life. In both cases, as new data comes in, the model becomes better and starts moving away from the fleet prior and converges to the actual battery. As expected, for battery 3, when data added near the "knee" around $\Delta e = 1$ kWh causes significant improvement of the median posterior curve.

Finally, we also studied how our proposed Bayesian approach compares against a simple curve fitting of the data. Figure 6.7 shows the capacity data and models versus Δe. In green we show our Bayesian approach, while in blue we show a linear regression model. The solid lines are the 50-percentile and the dashed lines are the 95% prediction intervals. For each of the three batteries, we illustrate a first case in which only few data points are available (top) and a second case in which more data points are available (bottom). The difference between Fig. 6.7a and Fig. 6.7b is the proportion of unknown history. The difference between Fig. 6.7a and Fig. 6.7c is the fact that battery 3 is closer to the fleet average, while battery 1 sits on the upper tail of the distribution.

The main point of Fig. 6.7 is that our variational hybrid physics-informed neural networks captures the overall form of the capacity versus Δe curve through the fleet-wide prior as opposed to simply fitting the data. The simple regression models might even fit the observed data well (pass through the points); but they lack any information regarding overall behavior of these curves. The collected points might even fall into a straight line (or exhibit only small curvature); and therefore, give the impression that a regression model could be used here. However, linear regression fails to capture the sudden drop in capacity as battery ages (and chasing the degree of the polynomial that could fit the data can be a frustrating exercise). On the other hand, our proposed model has uncertainty bands that initially reflect the fleet-wide uncertainty. The model converges and successfully captures battery-specific uncertainty as more data becomes available.
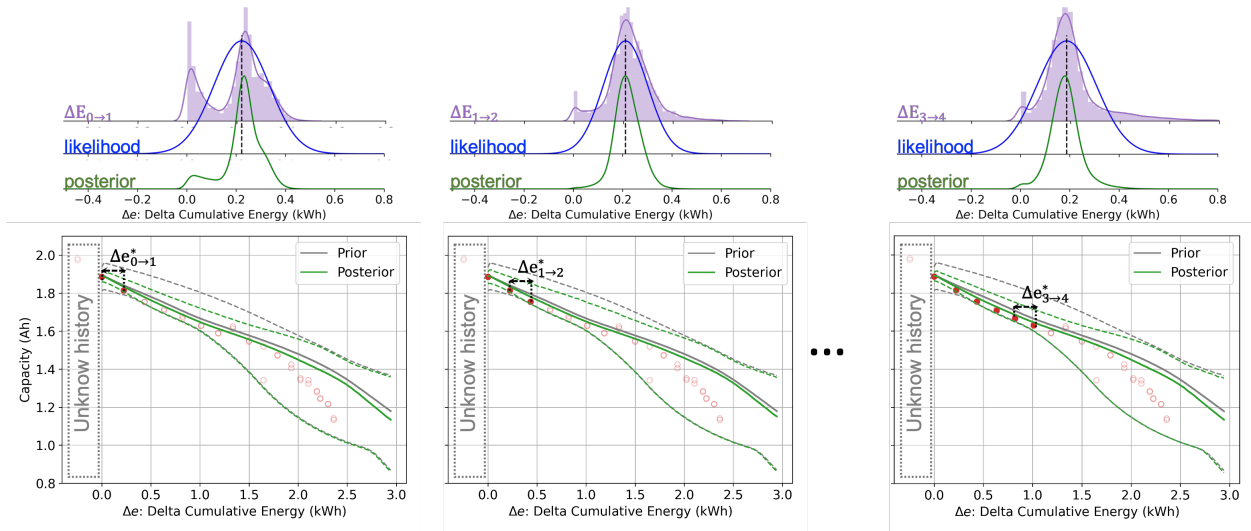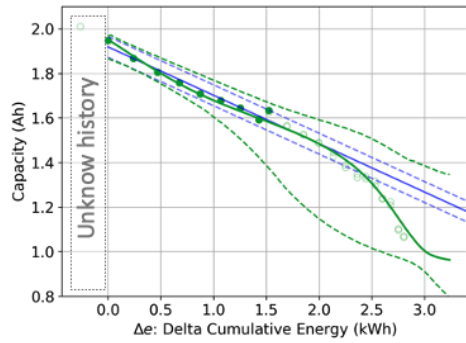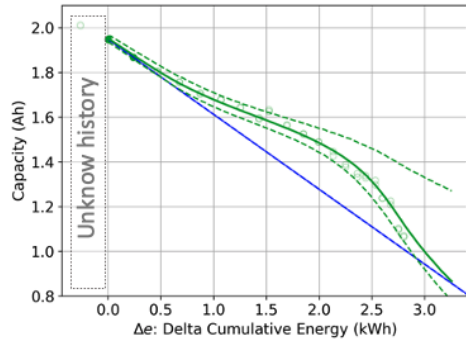
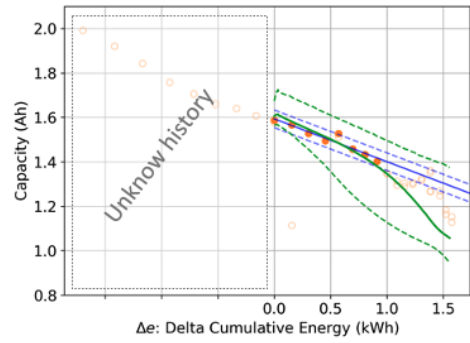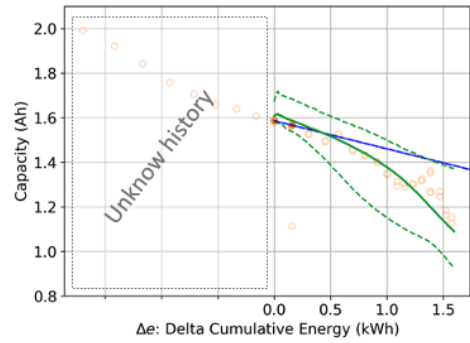(a) Battery 4, early life.
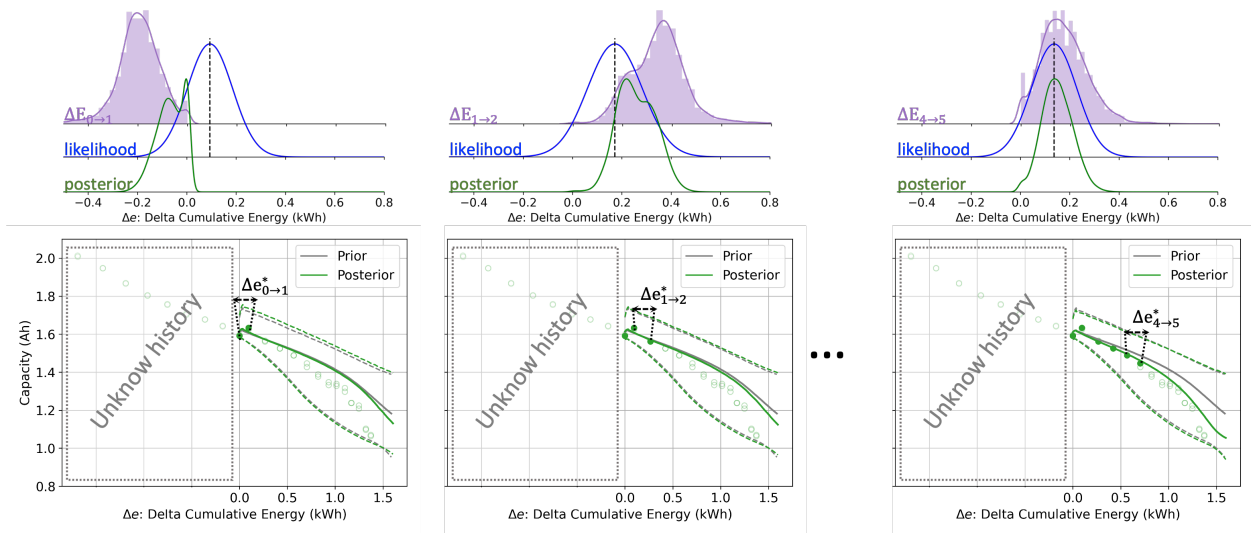
(b) Battery 3, mid life.

Figure 6.6: Sequential Bayesian update.

(a) Battery 3.

(b) Battery 2.
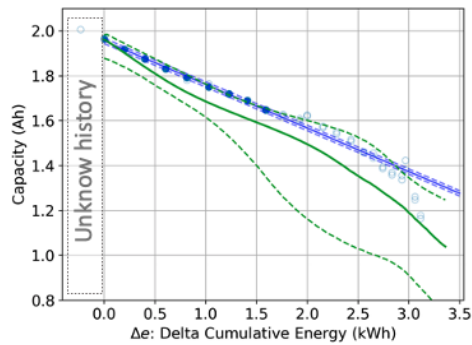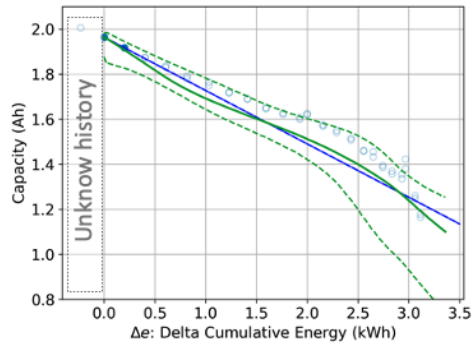
(c) Battery 1.

Figure 6.7: Comparison of battery-specific capacity vs Δe models.

CHAPTER 7: CLOSING REMARKS AND FUTURE WORK

Electric and hybrid systems are at the foundation of the current transportation transformation. Millions of people, in the short term, and possibly billions of people in the long term, are expected to use vehicles with some form of electric power train as a source of propulsion. As Li-ion and Li-Po batteries remain the most common way to provide energy to electric power trains, reliable tools to monitor the battery in real-time and perform diagnosis and prognosis are extremely important.

The main output of this research is a hybrid physics-informed neural networks model for monitoring the state of charge (SOC) of Lithium-ion batteries that can accurately estimate voltage discharge cycles, and track and forecast the aging of the battery over time.

The resultant prognosis model can accurate estimate the remaining useful capacity of lithium-ion batteries, able to forecast future discharge cycles, and deployable to be used on real fleets. Operators will be able to plug this prognosis model and have real-time state of charge (SOC) monitoring and battery remaining useful life estimation.

Our hybrid implementation of Li-ion battery prognosis is available in following the scientific publications [78], [79], [80]. These papers lay down fundamentals that can enable reliability assurance for missions, ensuring batteries will complete critical tasks. In addition, the model can be utilized to manage a fleet of batteries assets, optimizing battery usage as they age. At the foundation level (outside Li-ion batteries), the work on hybrid physics-informed neural networks was also published in venues such as [51, 75, 81].

Lessons Learned

In this research effort, we proposed a hybrid physics-informed neural network approach to Li-ion battery prognosis. The hybrid model combined known physics of discharge (such as Nernst and Butler-Volmer equations) with flexible machine learning for further uncertainty quantification and improved agreement with experimental data. We accounted for aging of the battery by modeling variation of q^{max} and R_0 as a function of cumulative energy draw from the battery, also using multi-layer perceptrons trained with variational inference. Our hybrid model framework has multiple advantages. The hybrid model was capable of predicting voltage discharge curves as well as end-of-discharge of batteries subject to constant or random loading conditions. Our choice to track (q^{max}, R_0) for each discharge curve enables tailoring of the model to become battery-specific. Therefore, the aging effect caused by a charge-discharge cycles can be captured by updating the parameter pair. The hybrid model can be easily tuned and interpreted, as most computations within the recurrent neural network cell are driven by the physics-based portion of the model. Yet, it takes advantage of the high-flexibility of multi-layer perceptrons to correct for missing physics and to account for uncertainty in some critical model parameters.

The aging model, built upon two variational multi-layer perceptrons used to estimate q^{max} and R_0 as a function of the cumulative energy draw from the battery, enables far-ahead prediction of voltage discharge curves by forecasting the values of q^{max} and R_0 far in the future. With the forecasted distributions of q^{max} and R_0, important operational outputs such as voltage drop and end of discharge can be predicted. In order to improve aging model forecasts, we proposed an ensemble model performing a weighted average of existing aging curves from older batteries. The fundamental hypothesis behind the approach is that the battery that is currently being monitored is a member of a fleet of similar batteries in a later aging state. Therefore, information from older batteries to forecast the capacity drop (through the proxies q^{max} and R_0) can be leveraged.

Uncertainty quantification using variational layers also successfully aid the aging forecast process. In all tested cases, predictions embraced the true discharge curves within their prediction intervals (see for example Figs. 4.4 and 4.5).

The approach is fully Bayesian as it uses fleet-wide capacity data to build prior distributions for battery-specific models and keeps updating these models with operational data. These are some of the major advantages of the proposed framework:

- *Quantification of model-form uncertainty:* the hybrid model expands reduced-order models with multi-layer perceptrons that compensate for model simplifications. As opposed to purely data-driven approaches, we demonstrated that our approach can successfully reduce the gap between predictions and observations with only a very reduced number of training points.

- *Quantification of battery-to-battery aging uncertainty:* the hybrid model implements a variational multi-layer perceptron to account for aging of batteries as a function of usage.

- *Efficient use of available data:* priors built using fleet-wide capacity data, thus reducing the need for constant discharge cycles are required for precise estimates of battery residual capacity. The model leverages existing data from a fleet of similar batteries to accurately predict its aging behavior, without the need to decommission the battery for ad-hoc testing.

- *Reliable forecasts:* the approach is able to monitor the predictive capability of the model through metrics such as KL-divergence. This allows the detection of diverging aging patterns (on a battery-specific basis), thus raising flags for ad-hoc testing to refine the models for those particular specimens.

Nevertheless, we also mention the following assumptions/disadvantages of the proposed approach:

- *First-principle models:* the approach builds on top of models that, at least, approximate the actual input-output relationship. Furthermore, the computational cost of the numerical implementation for these models need to be in the realm of the linear algebra found in neural networks; and there should be no impediment for backpropagation of gradients. This allows the implementation of neural network training, as well as fast and accurate predictions.

- *Fleet-wide information:* the approach relies on a number of existing batteries that have been previously monitored and their available capacity was recorded during operation. Otherwise, the fleet-wide prior models would be characterized by larger uncertainty, thus limiting the performance of the prediction.

- *Random discharge cycles:* the battery data-set utilized here contains many cases for which the used capacity is relatively close or approaches the available capacity. If the battery were repeatedly subject to lower loading conditions (for example, an UAV typically flying at speed and in environmental conditions well below the capability of its battery), the fleet-wide prior distribution would remain mostly unfiltered, and the predictions of future voltage discharge curve would be driven by the fleet information only. In addition, batteries could be recharged even if they were not fully or consistently depleted. This could introduce further errors in the modeling assumptions and contribute to premature divergence of models.

Future Research

As potential future research, the model could be tested under realistic loading profiles from small unmanned aerial vehicle flights and extended the physics-informed neural network model to more powertrain components as well as build a full hybrid powertrain model. A hybrid powertrain model would benefit not only end-of-discharge predictions, but also fault detection and isolation within the powertrain system.

In our implementation, most of the model is driven by first-principle equations, while neural network blocks describes the relationships between hidden variables that are yet not well understood or characterized by large inter-specimen variability. Therefore, these data-driven "blocks" compensate for discrepancies between model predictions and observed behavior. The approach is easily scalable to more complex models, for example, of the entire UAV power-train composed of DC motor, electronic speed controller, and battery. Further steps of this research could aim at relaxing some of the hypothesis this approach is based upon and extending the model to other power-train elements and testing it on realistic usage data.

Another venue that could be explored is the model architecture alternatives (number of layers, number of neurons, activations functions, and choice of optimizers) of the data-driven nodes. Subject to application complexity, one could pursue the optimization of the data-driven architecture alternatives with technologies such as neural architecture search [82–84].

LIST OF REFERENCES

[1] Matthew J Daigle and Chetan Shrikant Kulkarni. Electrochemistry-based battery modeling for prognostics. In *Annual Conference of the Prognostics and Health Management Society*, page 13 pages, New Orleans, USA, 9 2013. PHM Society.

[2] B Bole, C Kulkarni, and M Daigle. Randomized battery usage data set. *NASA AMES prognostics data repository*, 70, 2014.

[3] Christian Friedrich and Paul A Robertson. Hybrid-electric propulsion for aircraft. *Journal of Aircraft*, 52(1):176–189, 2015.

[4] Nateri Madavan, J Heidmann, C Bowman, P Kascak, A Jankovsky, and R Jansen. A NASA perspective on electric propulsion technologies for commercial aviation. In *Proceedings of the Workshop on Technology Roadmap for Large Electric Machines, Urbana-Champaign, IL, USA*, pages 5–6, 2016.

[5] Jimi Russell. NASA, U.S. industry aim to electrify commercial aviation, 10 2019.

[6] G Girishkumar, B McCloskey, Alan C Luntz, Sally Swanson, and W Wilcke. Lithium-air battery: promise and challenges. *The Journal of Physical Chemistry Letters*, 1(14):2193–2203, 2010.

[7] Robert F. Service, 5 2016.

[8] Jian Zhi, Alireza Zehtab Yazdi, Gayathri Valappil, Jessica Haime, and Pu Chen. Artificial solid electrolyte interphase for aqueous lithium energy storage systems. *Science Advances*, 3(9):e1701010, 2017.

[9] Ajay Misra. Energy conversion and storage requirements for hybrid electric aircraft. In *40th*

International Conference and Expo on Advanced Ceramics and Composites, page 9 pages. PHM Society, 1 2016.

[10] D. Eroglu, K. R. Zavadil, and K. G. Gallagher. Critical link between materials chemistry and cell-level design for high energy density and low cost lithium-sulfur transportation battery. *Journal of the Electrochemical Society*, 162(6), 1 2015.

[11] Zhaoyi Xu and Joseph Homer Saleh. Machine learning for reliability engineering and safety applications: Review of current status and future opportunities. *Reliability Engineering & System Safety*, 211:107530, 2021.

[12] Yigit A. Yucesan, Arinan Dourado, and Felipe A. C. Viana. A survey of modeling for prognosis and health management of industrial equipment. *Advanced Engineering Informatics*, 50:101404, 2021.

[13] Yoichi Takagishi, Takumi Yamanaka, and Tatsuya Yamaue. Machine learning approaches for designing mesoscale structure of li-ion battery electrodes. *Batteries*, 5(3):54, aug 2019.

[14] Kan Hatakeyama-Sato, Toshiki Tezuka, Yoshinori Nishikitani, Hiroyuki Nishide, and Kenichi Oyaizu. Synthesis of lithium-ion conducting polymers designed by machine learning-based prediction and screening. *Chemistry Letters*, 48(2):130–132, feb 2019.

[15] Zhuoyuan Zheng, Bo Chen, Yanwen Xu, Nathan Fritz, Yashraj Gurumukhi, John Cook, Mehmet N. Ates, Nenad Miljkovic, Paul V. Braun, and Pingfeng Wang. A Gaussian process-based crack pattern modeling approach for battery anode materials design. *Journal of Electrochemical Energy Conversion and Storage*, 18(1), may 2020.

[16] Arno Kwade, Wolfgang Haselrieder, Ruben Leithoff, Armin Modlinger, Franz Dietrich, and Klaus Droeder. Current status and challenges for automotive battery production technologies. *Nature Energy*, 3(4):290–300, apr 2018.

[17] Mona Faraji Niri, Kailong Liu, Geanina Apachitei, Luis Roman Ramirez, Michael Lain, Dhammika Widanage, and James Marco. Machine learning for optimised and clean Li-ion battery manufacturing: Revealing the dependency between electrode and cell characteristics. *Journal of Cleaner Production*, 324:129272, 2021.

[18] Artem Turetskyy, Jacob Wessel, Christoph Herrmann, and Sebastian Thiede. Battery production design using multi-output machine learning models. *Energy Storage Materials*, 38:93–112, 2021.

[19] Bhaskar Saha, Scott Poll, Kai Goebel, and Jon Christophersen. An integrated approach to battery health monitoring using bayesian regression and state estimation. In *2007 IEEE Autotestcon*, pages 646–653. IEEE, 2007.

[20] Fan Xu, Fangfang Yang, Zicheng Fei, Zhelin Huang, and Kwok-Leung Tsui. Life prediction of lithium-ion batteries based on stacked denoising autoencoders. *Reliability Engineering & System Safety*, 208:107396, 2021.

[21] Ting Tang and Huimei Yuan. A hybrid approach based on decomposition algorithm and neural network for remaining useful life prediction of lithium-ion battery. *Reliability Engineering & System Safety*, 217:108082, 2022.

[22] Xiaoyu Li, Zhenpo Wang, and Jinying Yan. Prognostic health condition for lithium battery using the partial incremental capacity and Gaussian process regression. *Journal of Power Sources*, 421:56–67, 2019.

[23] Linxia Liao and Felix Köttig. A hybrid framework combining data-driven and model-based methods for system remaining useful life prediction. *Applied Soft Computing*, 44:191–199, 2016.

[24] Cheng Fan, Fu Xiao, Chengchu Yan, Chengliang Liu, Zhengdao Li, and Jiayuan Wang. A

novel methodology to explain and evaluate data-driven building energy performance models based on interpretable machine learning. *Applied Energy*, 235:1551–1560, 2019.

[25] David Gunning and David W Aha. DARPA's explainable artificial intelligence program. *AI Magazine*, 40(2):44–58, 2019.

[26] Xun Jia, Lei Ren, and Jing Cai. Clinical implementation of ai technologies will require interpretable ai models. *Medical physics*, 47(1):1–4, 2020.

[27] Wojciech Samek, Grégoire Montavon, Andrea Vedaldi, Lars Kai Hansen, and Klaus-Robert Müller. *Explainable AI: interpreting, explaining and visualizing deep learning*, volume 11700. Springer Nature, 2019.

[28] Gregory L Plett. Extended Kalman filtering for battery management systems of lipb-based hev battery packs: Part 2. modeling and identification. *Journal of Power Sources*, 134(2):262–276, 2004.

[29] Guijun Ma, Yong Zhang, Cheng Cheng, Beitong Zhou, Pengchao Hu, and Ye Yuan. Remaining useful life prediction of lithium-ion batteries based on false nearest neighbors and a hybrid neural network. *Applied Energy*, 253:113626, 2019.

[30] Shuangqi Li, Hongwen He, Chang Su, and Pengfei Zhao. Data driven battery modeling and management method with aging phenomenon considered. *Applied Energy*, 275:115340, 2020.

[31] Matthieu Dubarry, Cyril Truchot, and Bor Yann Liaw. Synthesize battery degradation modes via a diagnostic and prognostic model. *Journal of Power Sources*, 219:204–216, 2012.

[32] Homero Valladares, Tianyi Li, Likun Zhu, Hazim El-Mounayri, Ahmed M. Hashem, Ashraf E. Abdel-Ghany, and Andres Tovar. Gaussian process-based prognostics of lithium-

ion batteries and design optimization of cathode active materials. *Journal of Power Sources*, 528:231026, 2022.

[33] Xiaopeng Tang, Changfu Zou, Ke Yao, Jingyi Lu, Yongxiao Xia, and Furong Gao. Aging trajectory prediction for lithium-ion batteries via model migration and Bayesian Monte Carlo method. *Applied Energy*, 254:113591, 2019.

[34] Sheng Shen, Mohammadkazem Sadoughi, Meng Li, Zhengdao Wang, and Chao Hu. Deep convolutional neural networks with ensemble learning and transfer learning for capacity estimation of lithium-ion batteries. *Applied Energy*, 260:114296, 2020.

[35] Yujie Cheng, Dengwei Song, Zhenya Wang, Chen Lu, and Noureddine Zerhouni. An ensemble prognostic method for lithium-ion battery capacity estimation based on time-varying weight allocation. *Applied Energy*, 266:114817, 2020.

[36] Bernd Möller and Michael Beer. Engineering computation under uncertainty – capabilities of non-traditional models. *Computers & Structures*, 86(10):1024 – 1041, 2008.

[37] Matthew E. Riley and Ramana V. Grandhi. Quantification of model-form and predictive uncertainty for multi-physics simulation. *Computers & Structures*, 89(11):1206 – 1213, 2011.

[38] M. C. Kennedy and Anthony O'Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3):425–464, 2001.

[39] John McFarland, Sankaran Mahadevan, Vicente Romero, and Laura Swiler. Calibration and uncertainty analysis for computer simulations with multivariate output. *AIAA Journal*, 46(5):1253–1265, 2008.

[40] Benjamin Peherstorfer, Karen Willcox, and Max Gunzburger. Survey of multifidelity methods in uncertainty propagation, inference, and optimization. *SIAM Review*, 60(3):550–591, 2018.

[41] Dongkun Zhang, Lu Lu, Ling Guo, and George Em Karniadakis. Quantifying total uncertainty in physics-informed neural networks for solving forward and inverse stochastic problems. *Journal of Computational Physics*, 397:108850, 2019.

[42] M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686 – 707, 2019.

[43] Maziar Raissi. Deep hidden physics models: deep learning of nonlinear partial differential equations. *Journal of Machine Learning Research*, 19(25):1–24, 2018.

[44] Jin-Long Wu, Heng Xiao, and Eric Paterson. Physics-informed machine learning approach for augmenting turbulence models: A comprehensive framework. *Physical Review Fluids*, 3(7):074602, 2018.

[45] Jan S. Hesthaven and Stefano Ubbiali. Non-intrusive reduced order modeling of nonlinear problems using neural networks. *Journal of Computational Physics*, 363:55–78, 2018.

[46] Renee Swischuk, Laura Mainini, Benjamin Peherstorfer, and Karen Willcox. Projection-based model reduction: Formulations for physics-based machine learning. *Computers & Fluids*, 179:704–717, 2019.

[47] E. Samaniego, C. Anitescu, S. Goswami, V.M. Nguyen-Thanh, H. Guo, K. Hamdia, X. Zhuang, and T. Rabczuk. An energy approach to the solution of partial differential equations in computational mechanics via machine learning: Concepts, implementation and applications. *Computer Methods in Applied Mechanics and Engineering*, 362:112790, 2020.

[48] Somdatta Goswami, Cosmin Anitescu, Souvik Chakraborty, and Timon Rabczuk. Transfer learning enhanced physics informed neural network for phase-field modeling of fracture. *Theoretical and Applied Fracture Mechanics*, 106, april 2020.

[49] Andrew K. S. Jardine, Daming Lin, and Dragan Banjevic. A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mechanical Systems and Signal Processing*, 20(7):1483–1510, 2006.

[50] Deepak K Karthikeyan, Godfrey Sikha, and Ralph E White. Thermodynamic model development for lithium intercalation electrodes. *Journal of Power Sources*, 185(2):1398–1407, 2008.

[51] Renato G. Nascimento and Felipe A. C. Viana. Cumulative damage modeling with recurrent neural networks. *AIAA Journal*, 58(12):5459–5471, 2020.

[52] Arinan Dourado and Felipe A. C. Viana. Physics-informed neural networks for missing physics estimation in cumulative damage models: a case study in corrosion fatigue. *ASME Journal of Computing and Information Science in Engineering*, 20(6):061007 (10 pages), 2020.

[53] Yigit A. Yucesan and Felipe A. C. Viana. A physics-informed neural network for wind turbine main bearing fatigue. *International Journal of Prognostics and Health Management*, 11(1):17 pages, 2020.

[54] Michael I Jordan, Zoubin Ghahramani, Tommi S Jaakkola, and Lawrence K Saul. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233, 1999.

[55] Gregory L. Plett. Extended Kalman filtering for battery management systems of LiPB-based HEV battery packs: Part 3. State and parameter estimation. *Journal of Power Sources*, 134(2):277–292, 2004.

[56] Chao Hu, Byeng D. Youn, and Jaesik Chung. A multiscale framework with extended Kalman

filter for lithium-ion battery SOC and capacity estimation. *Applied Energy*, 92:694–704, 2012.

[57] Nikolaos Wassiliadis, Jörn Adermann, Alexander Frericks, Mikhail Pak, Christoph Reiter, Boris Lohmann, and Markus Lienkamp. Revisiting the dual extended Kalman filter for battery state-of-charge and state-of-health estimation: A use-case life cycle analysis. *Journal of Energy Storage*, 19:73–87, 2018.

[58] Brian Bole, Chetan S Kulkarni, and Matthew Daigle. Adaptation of an electrochemistry-based Li-ion battery model to account for deterioration observed under randomized use. In *Annual Conference of the Prognostics and Health Management Society*, page 9 pages. PHM Society, 9 2014.

[59] Felipe A. C. Viana and Arun K. Subramaniyan. A survey of Bayesian calibration and physics-informed neural networks in scientific modeling. *Archives of Computational Methods in Engineering*, 28(5):3801–3830, Feb 2021.

[60] Barak A. Pearlmutter. Learning state space trajectories in recurrent neural networks. *Neural Computation*, 1(2):263–269, 1989.

[61] Alex Aussem. Dynamical recurrent neural networks towards prediction and modeling of dynamical systems. *Neurocomputing*, 28(1-3):207–232, 1999.

[62] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.

[63] David Barber. *Bayesian Reasoning and Machine Learning*. Cambridge University Press, 2012.

[64] C. Teubert, M. Corbetta, C. Kulkarni, and M. Daigle. NASA's prognostics models python package, 2020.

[65] B. Saha and K. Goebel. Battery data set. NASA Ames Research Center, Retrieved 11 May 2020, 2007.

[66] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[67] Ron Kohavi. A study of cross-validation and bootstrap for accuracy estimation and model selection. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, IJCAI'95, pages 1137–1143. Morgan Kaufmann Publishers Inc., 1995.

[68] Felipe A. C. Viana, Raphael T. Haftka, and Valder Steffen. Multiple surrogates: how cross-validation errors can help us to obtain the best predictor. *Structural and Multidisciplinary Optimization*, 39(4):439–457, 1 2009.

[69] Sean Wallis. Binomial confidence intervals and contingency tests: mathematical fundamentals and the evaluation of alternative methods. *Journal of Quantitative Linguistics*, 20(3):178–208, 2013.

[70] Lawrence D. Brown, T. Tony Cai, and Anirban DasGupta. Interval estimation for a binomial proportion. *Statistical Science*, 16(2):101–117, 2001.

[71] Solomon Kullback. *Information theory and statistics*. Courier Corporation, 1997.

[72] Alex Graves. Practical variational inference for neural networks. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 2348–2356. Curran Associates, Inc., 2011.

[73] Diederik Kingma and Max Welling. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*, 2014.

[74] Dustin Tran, Mike Dusenberry, Mark van der Wilk, and Danijar Hafner. Bayesian layers: A module for neural network uncertainty. In H. Wallach, H. Larochelle, A. Beygelzimer,

F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 14660–14672. Curran Associates, Inc., 2019.

[75] Renato G. Nascimento, Kajetan Fricke, and Felipe A. C. Viana. A tutorial on solving ordinary differential equations using Python and hybrid physics-informed neural networks. *Engineering Applications of Artificial Intelligence*, 96:103996, 2020.

[76] Dawn An, Joo-Ho Choi, and Nam Ho Kim. Prognostics 101: a tutorial for particle filter-based prognostics algorithm using Matlab. *Reliability Engineering & System Safety*, 115:161–169, 2013.

[77] Christian P. Robert, Víctor Elvira, Nick Tawn, and Changye Wu. Accelerating mcmc algorithms. *WIREs Computational Statistics*, 10(5):e1435, 2018.

[78] Renato G. Nascimento, Felipe A. C. Viana, Matteo Corbetta, and Chetan S. Kulkarni. Usage-based lifing of Lithium-ion battery with hybrid physics-informed neural networks. In *AIAA Aviation 2021 Forum*, Virtual Event, 8 2021. AIAA.

[79] Renato G. Nascimento, Matteo Corbetta, Chetan S. Kulkarni, and Felipe A. C. Viana. Hybrid physics-informed neural networks for lithium-ion battery modeling and prognosis. *Journal of Power Sources*, 513:230526, 2021.

[80] Renato G. Nascimento, Matteo Corbetta, Chetan S. Kulkarni, and Felipe A. C. Viana. Li-ion battery aging with hybrid physics-informed neural networks and fleet-wide data. In *2020 Annual Conference of the PHM Society*, Virtual Event, Nov-Dec 2021. PHM Society.

[81] Felipe A. C. Viana, Renato G. Nascimento, Arinan Dourado, and Yigit A. Yucesan. Estimating model inadequacy in ordinary differential equations with physics-informed neural networks. *Computers and Structures*, 245:106458, 2021.

[82] Kirthevasan Kandasamy, Willie Neiswanger, Jeff Schneider, Barnabas Poczos, and Eric P. Xing. Neural architecture search with Bayesian optimisation and optimal transport. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 2016–2025. Curran Associates, Inc., 2018.

[83] Chenxi Liu, Barret Zoph, Maxim Neumann, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, and Kevin Murphy. Progressive neural architecture search. In *The European Conference on Computer Vision (ECCV)*, Munich, Germany, September 2018. Computer Vision Foundation.

[84] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture search: a survey. *Journal of Machine Learning Research*, 20(55):1–21, 2019.