Robust Multi-Modal Hamilton-Jacobi Reachability Prognostics: An Application to Battery Health Management

Khoury Boutrous[1], Abdel Rahman El Khatib[1,2], Ghaleb Hoblos[2], Kokou Langueh[1], Eric Duviella[1], and Jacques Boonaert[1]

[1] *IMT Nord Europe, Institut Mines-Télécom, Univ. Lille, Centre for Education, Research and Innovation in Digital Systems (CERI Digital Systems), F-59500 Douai, France.*

[2] *University of Rouen Normandie, École Supérieure d'Ingénieurs en Génie Électrique (ESIGELEC), Normandie University, Institut de Recherche en Systèmes Électroniques Embarqués (IRSEEM), UR4353, F-76000 Rouen, France.*

ABSTRACT

This paper presents a novel and robust prognostics framework for battery health management based on Hamilton–Jacobi reachability analysis. A two-dimensional degradation state space is constructed from the state of health, obtained from discharge-capacity measurements, and a normalized impedance feature extracted from electrochemical impedance spectroscopy. Within this joint state space, a failure region is defined to capture both capacity fade and impedance growth under a unified EOL criterion. The Hamilton–Jacobi partial differential equation is then solved backward in a minimum-time-to-reach setting to generate state-dependent remaining useful life maps. To account for battery-to-battery variability, uncertainty in the degradation dynamics is estimated empirically from experimental data and incorporated through nominal, worst-case, and best-case drift scenarios, thereby yielding corresponding remaining useful life predictions. The resulting maps are computed offline and interpolated online, making the framework computationally efficient in deployment. Validation on the NASA battery dataset shows that the proposed approach delivers physically interpretable remaining useful life estimates together with informative uncertainty bounds that successfully encapsulate the true remaining useful life trajectories.

1. INTRODUCTION

Predicting, at the present time, the time remaining until a future event occurs is inherently uncertain and must therefore be accompanied by a mathematically quantified representation of the unavoidable degree of doubt associated with such estimates. In this sense, uncertainty quantification (UQ) is not optional in remaining useful life (RUL) prediction, but a fundamental requirement (Boutrous, Thuillier, Jha, Puig, & Theilliol, 2023). Although end users may often prefer point estimates for simplicity, particularly in consumer-oriented applications, such estimates alone may provide limited information if they are not accompanied by a measure of confidence. Incorporating uncertainty bounds into RUL predictions provides more informative decision support (Lin, Yan, & Zio, 2025). It enables engineers to assess the reliability of prognostic outputs, offer valuable feedback for improving prediction algorithms, sensing strategies, and data collection procedures, ultimately leading to more trustworthy and effective prognostic systems. Indeed, the absence of UQ in many state-of-the-art prognostic methodologies remains a major barrier to their adoption, particularly for black-box deep learning approaches (Basora, Viens, Arias Chao, & Olive, 2025). The importance of UQ in prognostics has been widely recognized and carefully examined in the literature (Ding, Qin, Zhang, & Lin, 2024; Nemani et al., 2023). A wide range of methods has been proposed in the literature to address both aleatoric and epistemic uncertainties, each with its own advantages and limitations.

For example, Kalman-filter-based methods, e.g. (Nunes et al., 2023), are attractive for real-time prognostics because of their computational efficiency and their ability to propagate state estimates together with associated covariance information. However, their uncertainty representation is typically tied to specific modeling assumptions, most notably approximate linearity and Gaussian noise, which can limit their ability to capture complex degradation dynamics and nonstandard uncertainty structures. Particle filters are more flexible in this regard, as they can represent nonlinear and non-Gaussian state distributions and therefore provide a richer description of predictive uncertainty. This added expressiveness, however, often comes at the expense of significantly higher computational cost in online applications (Jouin, Gouriveau, &

Hissel, 2016). More broadly, many classical UQ methods in prognostics remain strongly dependent on distributional assumptions, such as Gaussianity, that may not faithfully represent the true aleatoric and epistemic uncertainties affecting degradation evolution and RUL prediction.

In deep learning–based RUL prediction, methods such as Bayesian neural networks, Monte Carlo dropout, deep ensembles, and conformal prediction have also been used to estimate predictive uncertainty. However, these approaches can be computationally expensive, sometimes poorly calibrated, and in many cases rely on heuristic approximations rather than fully probabilistic foundations (Nemani et al., 2023).

Recent studies have also explored set-based bounding and reachability analysis (Al-Mohamad, Puig, & Hoblos, 2023; Boutrous, Puig, & Nejjari, 2022). Although these methods offer guaranteed uncertainty coverage and are appealing for safety-critical applications, they often yield overly conservative results, leading to confidence bounds that may be excessively wide and insufficiently informative in practice. UQ in battery prognostics has also been comprehensively reviewed in recent literature (Helen, Huan, Paulson, et al., 2024), further highlighting both the growing importance of reliable uncertainty estimates and the need for methods that remain informative, interpretable, and practically useful.

Multimodal prognostics offers clear advantages by improving robustness and accuracy when a single modality is insufficient to capture important degradation characteristics (Jose et al., 2025). However, it also introduces challenges related to heterogeneous data sources, differing sampling rates, noise levels, and the complexity of feature alignment and fusion, making such models harder to train and motivating the need for UQ. UQ in multimodal environments is particularly difficult when one modality dominates the prediction of the health index particularly in weighted black box scenarios such as in deep learning, which can lead to biased or poorly calibrated results. Hence, a framework that can overcome these limitations and deliver a reliable prognostic solution is needed.

In this paper, the prognostics of a fleet of batteries with similar chemistry is investigated by representing fleet variability in a multidimensional state space. The end-of-life (EOL) condition is formulated as a joint failure region in this state space, accounting for the contribution of all considered modalities. Based on this geometric failure boundary, a backward reachable HJ analysis is carried out to compute RUL value functions over the entire state space through a minimum-time-to-reach formulation, resulting in an offline RUL map, enabling an efficient computational time to an otherwise computationally time demanding endeavour. Although state-level fusion is not new in model-based prognostics, the proposed framework distinguishes itself by linking the fused degradation state to a joint geometric EOL boundary and an HJ-based reachable formulation. This yields an interpretable and robust approach to UQ, while remaining resilient to heterogeneous, dominant, missing, or degraded modalities. Furthermore, it avoids restrictive Gaussian assumptions which may be unjustifiable in data-scarce settings. The rest of the paper is organised as follows: The next section introduces the concept of reachability analysis, the HJ framework, and its extension to RUL prediction. A brief description of the dataset and the Preprocessing steps applied is then provided in Section 3. Section 4 presents the application of the HJB reachability framework to battery prognostics, while Section 5 discusses the results and its discussions. Finally, the conclusions are drawn in Section 6.

2. HAMILTON–JACOBI REACHABILITY FOR PROGNOSTICS

In this section, reachability analysis, and in particular HJ reachability analysis, is introduced. Building on this framework, RUL prediction is formulated through backward reachable value functions defined over the state space, which in turn give rise to RUL maps.

2.1. Reachability Analysis for Prognostics

Reachability analysis, a well-established tool in robotics, control, and safety verification, provides a rigorous mathematical framework for characterizing the evolution of dynamical systems under uncertainty. It encompasses the dual concepts of forward and backward reachability. In prognostics, these concepts can be leveraged to describe the evolution of degradation states and to assess their relationship to failure in the presence of uncertainty. Formally, the forward and backward reachable sets adapted to prognostics can be defined as follows:

Consider the dynamical system

$$\dot{x}(\tau) = f\big(x(\tau), d(\tau)\big), \qquad \tau \in [0, t], \tag{1}$$

where $x(\tau) \in \mathbb{R}^n$ is the degradation state vector and $d(\tau) \in \mathcal{D}$ represents admissible uncertainties. Let $\mathcal{X}_0 \subseteq \mathbb{R}^n$ denote the set of admissible initial states. Also the degradation states are assumed to be sufficiently informative so that future evolution can be approximated as Markovian in *Equation 1*. Then the following definitions hold:

Definition 1 (Forward Reachable Set (FRS)). The forward reachable set at time t from the initial set $\mathcal{X}_0$ is defined as the set of all states that can be attained by the system at time t from some initial condition $x_0 \in \mathcal{X}_0$, under the admissible degradation dynamics and uncertainties. Formally,

$$\mathcal{R}^+(t; \mathcal{X}_0) = \Big\{ x \in \mathbb{R}^n \;\Big|\; \exists x_0 \in \mathcal{X}_0,\ \exists d(\cdot) \in \mathcal{D}, \\ x = \xi(t; x_0, d(\cdot)) \Big\}. \tag{2}$$

where $\xi(t; x_0, d(\cdot))$ denotes the solution trajectory of the sys-

tem evaluated at time t, originating from the initial condition x_0 under the admissible uncertainty signal $d(\cdot)$.

Accordingly, $\mathcal{R}^+(t;\mathcal{X}_0)$ characterizes the set of all possible future degradation states that the system may occupy at time t under bounded uncertainty.

Definition 2 (Backward Reachable Set (BRS)). Given a target set $\mathcal{T} \subseteq \mathbb{R}^n$, the backward reachable set at time t is defined as the set of all initial states from which the system can evolve into the target set within the time horizon t, under the admissible degradation dynamics and uncertainties. Formally,

$$\mathcal{R}^-(t;\mathcal{T}) = \left\{ x_0 \in \mathbb{R}^n \quad \begin{array}{l} \exists d(\cdot) \in \mathcal{D}, \\ \text{such that } \xi(t; x_0, d(\cdot)) \in \mathcal{T} \end{array} \right\}. \tag{3}$$

In prognostics, when the target set $\mathcal{T}$ is selected as the EOL or failure region, the backward reachable set represents all degradation states from which failure can be reached within the time horizon t. This makes backward reachability particularly relevant to RUL prediction. By selecting an appropriate target set corresponding to the EOL region, a minimum-time backward reachability formulation provides a natural characterization of the RUL. Thus for a given target set $\mathcal{T} \subseteq \mathbb{R}^n$, the minimum-time-to-reach function is defined as

$$T(x_0) = \inf_{d(\cdot)\in\mathcal{D}} \{t_e \geq 0 \mid \xi(t_e; x_0, d(\cdot)) \in \mathcal{T}\}, \tag{4}$$

where $\xi(t_e; x_0, d(\cdot))$ denotes the state trajectory of the system starting from x_0 under the admissible uncertainty realization $d(\cdot)$ and t_e, the elapsed time from the initial state x_0 at which the state tube reached the target set $\mathcal{T}$. If the target set cannot be reached from x_0, then $T(x_0) = +\infty$. The minimum-time-to-reach function is directly related to the backward reachable set through

$$x_0 \in \mathcal{R}^-(t;\mathcal{T}) \iff T(x_0) \leq t. \tag{5}$$

2.2. The Hamilton-Jacobi BRS formulation

Several approaches have been proposed for the computation of backward reachable sets. Among these, the HJ framework is especially appealing because it can accommodate nonlinear dynamics, explicitly incorporate control and disturbance inputs, and represent reachable sets of arbitrary geometry (Bansal, Chen, Herbert, & Tomlin, 2017). Within this framework, the target set is encoded implicitly by means of a level-set function, and the associated HJ partial differential equation (PDE) is solved on a discretized state-space grid to obtain the corresponding value function. Its principal limitation, however, lies in the fact that the computational complexity increases exponentially with the dimension of the state space. Nevertheless, in the present setting this does not compromise practical implementation, since the RUL maps are computed offline and only interpolated during online deployment, thereby shifting the computational burden away from real-time operation and enabling efficient prognostic inference.

Therefore given a target set $\mathcal{T} \subseteq \mathbb{R}^n$. A level-set function $L : \mathbb{R}^n \to \mathbb{R}$ represents the set $\mathcal{T}$ implicitly through its zero sublevel set:

$$\mathcal{T} = \{x \in \mathbb{R}^n \mid L(x) \leq 0\}. \tag{6}$$

Accordingly, the boundary of the target set can be characterized by the zero level set

$$\partial\mathcal{T} = \{x \in \mathbb{R}^n \mid L(x) = 0\}, \tag{7}$$

with $L(x) < 0$ indicating that the state lies inside the target set, and $L(x) > 0$ indicating that it lies outside. In the prognostics setting considered here, the target set corresponds to the EOL region as stated previously, that is, the subset of the state space for which the system is considered to have reached failure. Consequently, the value function provides a natural way to characterize how close a given degradation state is to the EOL boundary and, in turn, supports the formulation of RUL predictions.

To compute the backward reachable set, the level-set representation of the target set is evolved according to the associated HJ Partial differential equation to obtain the corresponding value function. The resulting value function $V(x,t)$, at initial time $V(x,0) = L(x)$ satisfies

$$\frac{\partial V(x,t)}{\partial t} + \mathrm{opt}_{d\in D} \{\nabla V(x,t) \cdot f(x,d)\} = 0, \tag{8}$$

with terminal condition

$$V(x,0) = L(x). \tag{9}$$

Here, opt denotes either $\min$ or $\max$, depending on how the reachability problem is formulated and how uncertainty is modeled.

The backward reachable set at time t is then given by the zero sublevel set of the value function:

$$\mathcal{R}^-(t;\mathcal{T}) = \{x \in \mathbb{R}^n \mid V(x,t) \leq 0\}. \tag{10}$$

Thus, states satisfying $V(x,t) \leq 0$ can reach the target set within the prescribed horizon, whereas states for which $V(x,t) > 0$ cannot.

More generally, the backward reachable tube over the time horizon $[0,T]$ is defined as

$$\mathcal{R}^-_{\mathrm{tube}}(T) = \bigcup_{\tau\in[0,T]} \mathcal{R}^-(\tau;T) \tag{11}$$

This characterization is consistent with the definition of the backward reachable set given earlier in (3).

Although the above time-dependent HJ framework provides a general characterization of backward reachability and can, in principle, be used to define RUL, in this paper, a stationary minimum-time formulation is employed instead. In this context, stationary does not mean that time is neglected, rather it means that the value function is computed directly as the minimum time required for each state to reach the EOL set, without explicitly solving a time-dependent PDE.

2.3. HJ Minimum-time-to-reach interpretation for RUL

The HJ framework adopted in this paper yields a state-dependent RUL map by solving a stationary minimum-time-to-reach problem. Let $\mathcal{T}_{\mathrm{EOL}} \subseteq \mathbb{R}^n$ denote the EOL target set, and let $f^\ell : \mathbb{R}^n \to \mathbb{R}^n$ denote the vector field governing the state evolution under scenario ℓ. For any initial state $x_0 \in \mathbb{R}^n$, the RUL under scenario ℓ is defined as the minimum time required for the corresponding trajectory to reach the EOL set:

$$\mathrm{RUL}^\ell(x_0) = \inf\left\{\tau \geq 0 \;\middle|\; x^\ell(\tau; x_0) \in \mathcal{T}_{\mathrm{EOL}}\right\}, \qquad (12)$$

where $x^\ell(\tau; x_0)$ denotes the trajectory initialized at x_0 and governed by

$$\dot{x} = f^\ell(x). \qquad (13)$$

The associated value function $V^\ell(x)$ therefore directly represents the state-dependent RUL:

$$V^\ell(x) = \mathrm{RUL}^\ell(x). \qquad (14)$$

It satisfies the stationary HJ equation

$$\nabla V^\ell(x) \cdot f^\ell(x) = -1, \qquad x \in \Omega \setminus \mathcal{T}_{\mathrm{EOL}}, \qquad (15)$$

subject to the boundary condition

$$V^\ell(x) = 0, \qquad x \in \mathcal{T}_{\mathrm{EOL}}. \qquad (16)$$

If the EOL set is not reachable from a given state under scenario ℓ, then $V^\ell(x) = +\infty$. In practice, this is approximated numerically by truncating the solution over a finite computational horizon. A derivation of (15) is provided in Appendix A.

In this work, uncertainty is incorporated through three empirically constructed drift fields corresponding to nominal, worst-case, and best-case degradation scenarios. Drift field represents the local direction and rate of degradation evolution in the state space. Solving the stationary minimum-time problem for each drift field yields the corresponding RUL maps $V^{\mathrm{nom}}(x)$, $V^{\mathrm{wc}}(x)$, and $V^{\mathrm{bc}}(x)$, which provide nominal, pessimistic, and optimistic estimates of the RUL, respectively. Together, these maps produces a scenario-based characterization of prognostic uncertainty over the state space.

3. DATASET DESCRIPTION AND STATE-SPACE CONSTRUCTION

A brief introduction to the dataset is given. This is followed by a description of the preprocessing procedure and the feature extraction steps used to construct the multi-model state space. The section concludes with the estimation of the drift fields.

3.1. Selection of Battery Group

For brevity, a comprehensive description of the NASA battery dataset is not repeated here, interested readers are however referred to the original dataset documentation (Saha & Goebel, 2007). The dataset consists of controlled cycle-to-cycle experiments performed on 28 batteries, organized into 7 groups, all tested at room temperature. The recorded measurements include terminal current, voltage, and cell temperature during operation, as well as discharge capacity and Electrochemical Impedance Spectroscopy (EIS) measurements. In this paper, the state space is constructed using the state of health (SOH), derived from the discharge–capacity, together with a feature extracted from the electrochemical impedance measurements from Group6, whose cells are deemed to share similar chemistry.

3.2. Preprocessing and Feature Extraction

The monotonic decreasing transformation is applied to each battery's SOH sequence which is calculated from

$$s_k^{(j)} = \frac{Capacity_k}{C_{nom}}$$

Where, C_{nom} is the nominal capacity. Thus in order to preserve the physically meaningful degradation trend, the following procedure is administered. Let

$$s^{(j)} = \left[s_0^{(j)}, s_1^{(j)}, \ldots, s_{N_j-1}^{(j)}\right]$$

denote the SOH sequence of the j-th battery over N_j cycles. For each battery, the raw SOH sequence is first smoothed using a rolling-window average, producing a sequence $\tilde{s}^{(j)}$. Monotonicity is then enforced through the cumulative minimum operation

$$s_i^{(j)} = \min\left\{\tilde{s}_0^{(j)}, \tilde{s}_1^{(j)}, \ldots, \tilde{s}_i^{(j)}\right\}, \qquad i = 0, 1, \ldots, N_j - 1.$$

This guarantees that the resulting sequence $s^{(j)}$ is nonincreasing with cycle number. Thus the forced monotonicity imposes physically irreversibility assumptions associated with the battery aging. The transformed SOH trajectories for all batteries are shown in *Fig. 1*.

In addition to SOH, an impedance-based degradation feature is constructed from the charge-transfer resistance. For the j-

th battery, let

$$R_{ct}^{(j)} = \left[R_{ct,0}^{(j)}, R_{ct,1}^{(j)}, \ldots, R_{ct,N_j-1}^{(j)}\right]$$

denote the sequence of charge-transfer resistance measurements. The corresponding transformed feature is defined as

$$x_i^{(j)} = \log\left(\frac{R_{ct,i}^{(j)}}{R_{ct,0}^{(j)}}\right), \qquad i = 0, 1, \ldots, N_j - 1,$$

where $R_{ct,0}^{(j)}$ is the initial charge-transfer resistance of the j-th battery. This logarithmic normalization yields a dimensionless feature that measures the relative growth of internal resistance with respect to its initial value.

To reduce measurement noise, the raw impedance-derived sequence is first smoothed using a rolling-window average, producing $\tilde{x}^{(j)}$. Since resistance growth is expected to be nondecreasing during degradation, monotonicity is enforced through the cumulative maximum operation

$$\hat{x}_i^{(j)} = \max\left\{\tilde{x}_0^{(j)}, \tilde{x}_1^{(j)}, \ldots, \tilde{x}_i^{(j)}\right\}, \qquad i = 0, 1, \ldots, N_j - 1.$$

As a result, the processed feature $\hat{x}^{(j)}$ is monotonic increasing with cycle number. Together, the monotonic SOH sequence $m^{(j)}$ and the monotonic impedance feature $\hat{x}^{(j)}$ define the two-dimensional degradation state used in the multi-model state-space formulation:

$$z_i^{(j)} = \begin{bmatrix} s_i^{(j)} \\ \hat{x}_i^{(j)} \end{bmatrix}.$$

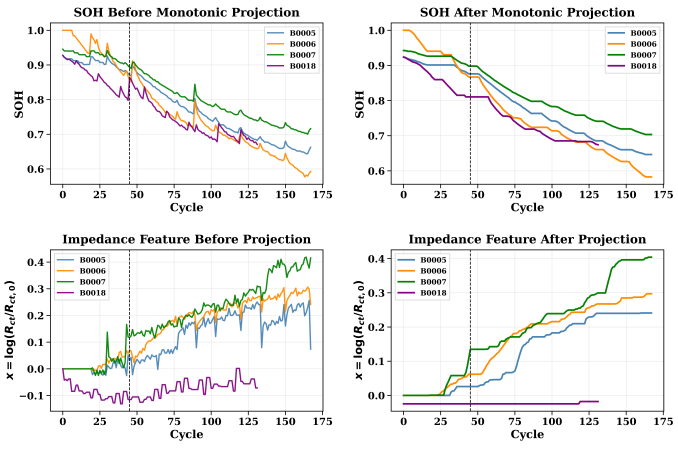

Figure 1. Monotonic decreasing and decreasing degradation health index of SOH and x from the impedance.(dashed line indicates the training/test partition)

These processed trajectories provide a physically consistent representation of battery degradation, with SOH capturing capacity fade and the impedance-based feature capturing resistance growth. In *Section 3.4* , these state trajectories are used to estimate the drift fields governing the degradation dynamics.

3.3. Assumptions Underlying the State-Space Construction

The proposed HJ prognostic framework related to this work relies on a two-dimensional degradation state constructed from measured battery data. The normalized discharge capacity, expressed through the SOH (s), is assumed to represent capacity fade, while the impedance-derived charge-transfer-resistance feature is assumed to capture internal resistance growth, regardless of any other exogenous factors. The logarithmic normalization

$$x_i^{(j)} = \log\left(\frac{R_{ct,i}^{(j)}}{R_{ct,0}^{(j)}}\right)$$

is used to make relative resistance growth more comparable across batteries. The state vector

$$z_i^{(j)} = \begin{bmatrix} s_i^{(j)} \\ \hat{x}_i^{(j)} \end{bmatrix}$$

therefore assumes that SOH and impedance provide complementary degradation information. The monotonic preprocessing further assumes that the dominant degradation trend is irreversible, with SOH decreasing and impedance increasing with cycle number. While this reduces measurement noise and local fluctuations, it may also remove short-term recovery or nonmonotonic effects. The method also assumes that the selected batteries are sufficiently similar in chemistry, testing protocol, and operating conditions to be represented in a common state space. Since cycle number is used as the independent variable, RUL is expressed in cycles, and local cycle-to-cycle increments are assumed to approximate the degradation drift. Consequently, a suitable dataset should provide repeated measurements of capacity-related and complementary degradation-sensitive features over enough cycles, under comparable operating conditions, with sufficient trajectory coverage of the relevant state-space regions. If these requirements are not met, the resulting drift estimates and HJ RUL maps should be interpreted with caution.

3.4. Drift field construction

Let $z = (s, \hat{x}) \in \mathbb{R}^2$ denote the degradation state, where s is the state of health and $\hat{x}$ is the impedance-based feature. From the measured battery trajectories, local increments $(\Delta s, \Delta \hat{x})$ i.e in equation 17) are computed in each partition of the state space to define an empirical drift field

$$f(z) = \begin{bmatrix} f_s(z) \\ f_{\hat{x}}(z) \end{bmatrix}.$$

Because only four batteries are available, the data are insufficient for a reliable probabilistic inference of drift uncertainty.

Using bootstrap resampling of the observed local increments within each state-space cell as discussed in the previous section, the, nominal, worst-case, and best-case drift fields, denoted by $f^{\mathrm{nom}}(z)$, $f^{\mathrm{wc}}(z)$, and $f^{\mathrm{bc}}(z)$, respectively are constructed as given below. The proposed drift-field construction represents cell-to-cell variability in the combined degradation state space by capturing the spread of observed local state increments across the battery trajectories. The algorithm is as given in Algorithm 1.

Algorithm 1 Construction of empirical drift fields from degradation trajectories

Require: Battery trajectories $\{(s_k^{(j)}, \hat{x}_k^{(j)})\}$, partition $\{\mathcal{B}_{qr}\}$, quantiles q_ℓ, q_h
Ensure: Drift fields $f^{\mathrm{nom}}(z)$, $f^{\mathrm{wc}}(z)$, $f^{\mathrm{bc}}(z)$
for $j = 1, \ldots, M$ **do**
  for $k = 0, \ldots, N_j - 2$ **do**
    Compute local increments $\Delta s_k^{(j)}$ and $\Delta \hat{x}_k^{(j)}$
    Compute midpoint state $\bar{z}_k^{(j)}$
  end for
end for
for each bin $\mathcal{B}_{qr}$ **do**
  Collect all increments whose midpoint lies in $\mathcal{B}_{qr}$
  Compute nominal drift using the componentwise median
  Compute worst-case drift using lower/upper bootstrap quantiles
  Compute best-case drift using upper/lower bootstrap quantiles
end for
Construct the piecewise fields $f^{\mathrm{nom}}(z)$, $f^{\mathrm{wc}}(z)$, and $f^{\mathrm{bc}}(z)$
Interpolate the fields if a continuous representation is required

Note that from Algorithm 1, M is the total number of batteries, N_j is the number of cycles recorded for the j-th battery, and $\mathcal{B}_{qr}$ denotes the (q, r)-th cell of the state-space partition, with q and r indexing the grid along the SOH and impedance directions, respectively. The local increments are defined by

$$\Delta s_k^{(j)} = s_{k+1}^{(j)} - s_k^{(j)}, \qquad \Delta \hat{x}_k^{(j)} = \hat{x}_{k+1}^{(j)} - \hat{x}_k^{(j)}, \tag{17}$$

and are associated with the midpoint state

$$\bar{z}_k^{(j)} = \left(\frac{s_k^{(j)} + s_{k+1}^{(j)}}{2}, \frac{\hat{x}_k^{(j)} + \hat{x}_{k+1}^{(j)}}{2} \right).$$

For each bin $\mathcal{B}_{qr}$, the nominal drift is defined as

$$f_{qr}^{\mathrm{nom}} = \begin{bmatrix} \mathrm{med}\{\Delta s_k^{(j)}\}_{\bar{z}_k^{(j)} \in \mathcal{B}_{qr}} \\ \mathrm{med}\{\Delta \hat{x}_k^{(j)}\}_{\bar{z}_k^{(j)} \in \mathcal{B}_{qr}} \end{bmatrix},$$

while the worst-case and best-case drifts are given by

$$f_{qr}^{\mathrm{wc}} = \begin{bmatrix} Q_{q_\ell}(\Delta s) \\ Q_{q_h}(\Delta \hat{x}) \end{bmatrix}, \qquad f_{qr}^{\mathrm{bc}} = \begin{bmatrix} Q_{q_h}(\Delta s) \\ Q_{q_\ell}(\Delta \hat{x}) \end{bmatrix}.$$

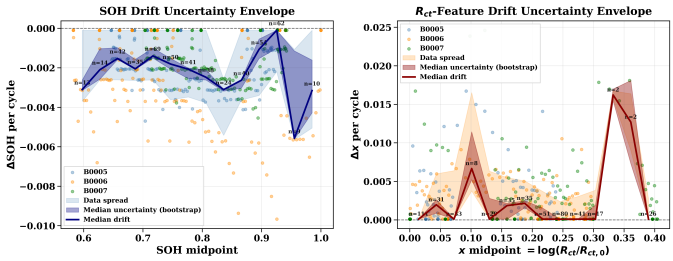

Figure 2. {Bootstrap drift uncertainty envelopes for SOH (left) and impedance feature (right). The solid line shows the nominal (median) per-cycle drift and the shaded band spans the worst-case to best-case quantile range, capturing battery-to-battery variability. SOH drift is negative (capacity fade) and impedance drift is positive (resistance growth), consistent with the enforced monotonicity.

The resulting cellwise drift estimates define three scenario-dependent degradation fields over the joint state space, corresponding to nominal, worst-case, and best-case evolution. These fields are interpreted dynamically as

$$\dot{z} = f^{\mathrm{nom}}(z), \qquad \dot{z} = f^{\mathrm{wc}}(z), \qquad \dot{z} = f^{\mathrm{bc}}(z),$$

and are used in the stationary HJ formulation to compute the corresponding RUL maps.

4. PROPOSED HJB-BASED MULTIMODAL PROGNOSTIC FRAMEWORK

The EOL region is defined as the superlevel target set

$$\mathcal{T}_{\mathrm{EOL}} = \{(s, x) \in \Omega : \alpha(1 - s) + \beta x \geq \gamma\},$$

where $\alpha, \beta > 0$ weight the contributions of capacity fade and resistance growth, respectively, and γ is the joint failure threshold. For a given drift field $f^\ell(z)$, with $\ell \in \{\mathrm{wc}, \mathrm{nom}, \mathrm{bc}\}$ denoting worst-case, nominal, and best-case scenarios, the RUL is defined as the minimum time required for the state trajectory to reach $\mathcal{T}_{\mathrm{EOL}}$. The corresponding value function $V^\ell(z)$ satisfies the stationary HJ equation

$$\nabla V^\ell(z) \cdot f^\ell(z) = -1, \qquad z \in \Omega \setminus \mathcal{T}_{\mathrm{EOL}},$$

subject to the boundary condition

$$V^\ell(z) = 0, \qquad z \in \mathcal{T}_{\mathrm{EOL}}.$$

the gradient of the value function is given by

$$\nabla V^\ell(z) = \begin{bmatrix} \frac{\partial V^\ell}{\partial s}(z) \\ \frac{\partial V^\ell}{\partial x}(z) \end{bmatrix},$$

while the degradation drift field is

$$f^\ell(z) = \begin{bmatrix} f_s^\ell(z) \\ f_x^\ell(z) \end{bmatrix}.$$

Hence, the HJ equation can be written explicitly as

$$\nabla V^{\ell}(z)^{\top} f^{\ell}(z) = -1 \tag{17}$$

This expression shows that the rate of change of the RUL map along the degradation trajectory is equal to -1, meaning that one unit of 'elapsed time' reduces the RUL by one unit. In particular, $\frac{\partial V^{\ell}}{\partial s}$ quantifies the sensitivity of the predicted RUL to changes in SOH, while $\frac{\partial V^{\ell}}{\partial x}$ quantifies the sensitivity of the predicted RUL to changes in the impedance-derived feature.

Numerically, this equation is solved on a structured (s, x) grid using an upwind discretization consistent with the monotone degradation directions $f_s^{\ell}(z) < 0$ and $f_x^{\ell}(z) > 0$. The resulting solutions $V^{\mathrm{wc}}, V^{\mathrm{nom}}, V^{\mathrm{bc}}$ define state-dependent pessimistic, nominal, and optimistic RUL maps, which are subsequently queried online by interpolation.

[1]

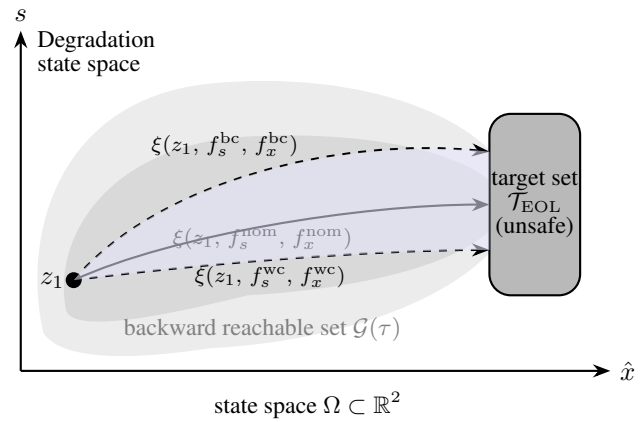

Figure 3. Schematic of the backward reachable set $\mathcal{G}(\tau)$ in the degradation state space $(s, \hat{x}) \in \Omega \subset \mathbb{R}^2$. Starting from z_1, the nominal, best-case, and worst-case drift realizations reach the unsafe target set $\mathcal{T}_{\mathrm{EOL}}$. These trajectories illustrate how uncertainty in the estimated drift field propagates into uncertainty in the predicted RUL.

5. RESULTS AND DISCUSSION

The weighting parameters are chosen as $\alpha = 0.8$, $\beta = 0.3$ and $\gamma = 0.2$, based on presumed designer domain knowledge and γ selected based on pre-caliberation for the EOL superlevel-set as shown in Fig.9. In this case more weight is placed on the SOH for the RUL prediction. This functional form is selected to reflect the monotonic degradation behavior of the two state variables, namely capacity fade and resistance growth. The HJ backward reachable sets are then computed under three degradation scenarios, yielding the corresponding pessimistic, nominal, and optimistic RUL maps used for evaluation.

$$\widehat{\mathrm{RUL}}_{\mathrm{nom}}(i) = V_n(s_i, x_i),$$
$$\widehat{\mathrm{RUL}}_{\mathrm{wc}}(i) = V_w(s_i, x_i),$$
$$\widehat{\mathrm{RUL}}_{\mathrm{bc}}(i) = V_b(s_i, x_i).$$

For model development, the first 45 cycles of each battery trajectory are used as the training dataset, while the remaining cycles are reserved for evaluation. For evaluation on measured trajectories $\{(s_k, x_k)\}$, the true EOL index is defined as the first entry into the same target set:

$$k_{\mathrm{EOL}} = \min\{k : \alpha(1 - s_k) + \beta x_k \geq \gamma\}. \tag{18}$$

The corresponding true RUL is then defined by

$$\mathrm{RUL}_{\mathrm{true}}(i) = \max\big(0, k_{\mathrm{EOL}} - i\big), \tag{19}$$

thereby ensuring consistency between the HJB-based prediction model and the empirical evaluation criterion.

Due to the limited amount of available data, the drift-field estimation is performed using the full length of the dataset, as illustrated in Fig. 2, which indicates coverage levels of 82% and 86% for SOH and R_{ct}, respectively. Figure 5 presents the various RUL maps obtained from the solution of the HJ problem to be interpolated for RUL values. Under the boundary condition defined in the state space, the figure illustrates the contour structure of the value function, which directly represents the RUL across the degradation domain. Furthermore, the uncertainty magnitude is also shown, providing an interpretable indication of the regions in the state space where prediction uncertainty is higher and thereby offering useful a posteriori insight into the confidence of the RUL estimates and possible directions for further refinement. The uncertainty is observed to be higher in the initial regions of the state space, predominantly due to the stronger influence of SOH on the degradation drift during the early stages, as also reflected in Fig. 2.

Fig. 4 illustrates the backward reachable tube on the nominal RUL map for the queried state $z = [0.8, 1.3]^{\top}$. The successive rings represent level sets of a family of value functions, propagated backward from the EOL boundary by a unit time-cycle slices from (17) (i.e. $c = 0, 1, 2, 3, \ldots$ cycles). Collectively, these level sets define the backward reachable tube. Since the queried state lies on the $c = 20$ contour, the corresponding predicted RUL is 20 cycles.

With the RUL maps evaluated offline, the corresponding bounded RUL estimates are queried online, as shown in Fig. 8. The purpose of this paper is primarily to introduce the proposed methodology and to provide a preliminary illustration of its capabilities, a more thorough comparison with existing methods is left for future work. However, in gen-

[1] This is a stationary minimum-time Hamilton–Jacobi equation rather than a time-dependent Hamilton–Jacobi PDE. In this formulation, the value function depends only on the state, so the explicit time derivative term $\partial V/\partial t$ is absent and the equation is posed directly over the state space through the Hamiltonian.

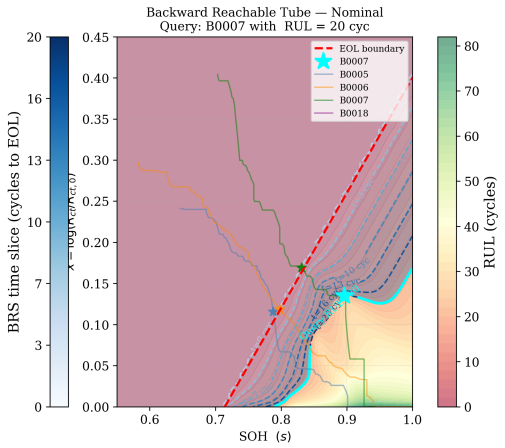

Figure 4. Illustration of the backward reachable tube on the nominal RUL map, propagated backward from the failure boundary to a queried state for predicting the RUL of battery B0007.

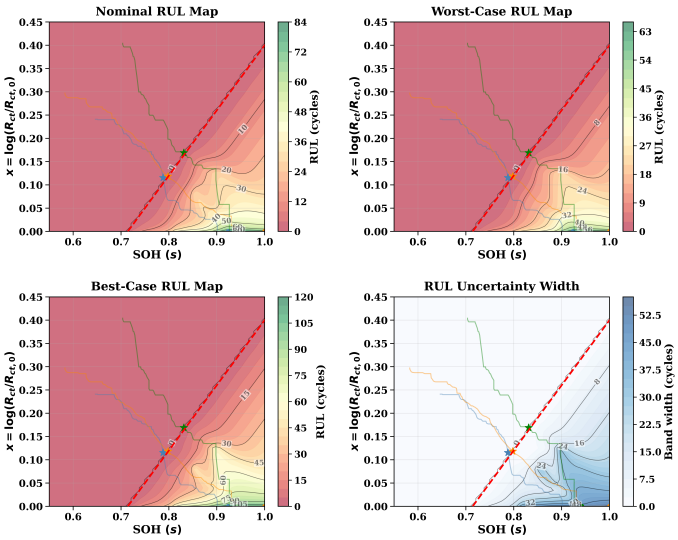

Figure 5. RUL maps in the multimodal HJ state space. From top left to bottom right: nominal condition, worst-case condition, best-case condition, and uncertainty band.

eral, Batteries B0005 and B0007 exhibit good coverage of the true RUL. In contrast, Battery B0018 shows poor coverage, with the predicted bounds failing to capture the true RUL over the evaluation horizon. This behavior mainly arises from the atypical degradation trend induced by the monotonic R_{ct} trajectory, as observed in Fig. 1, which causes the true RUL evolution to deviate significantly from the drift learned from the data. Battery B0006 also exhibits similar behavior. As evidenced by the drift uncertainty envelope in Fig. 2, the R_{ct} drift field is estimated from few datapoints in the high-x region traversed by B0006. Consequently, the learned drift and the resulting RUL map in that region are poorly supported, leading to a systematic underestimation of the RUL and a failure of the uncertainty bounds to capture the true RUL. This provides useful a posteriori information that can support further refinement and analysis. Indeed, the R_{ct} trajectory of B0018 and drift learning could be further improved through refined preprocessing; however, it is deliberately retained in its present form for illustrative purposes of the interpretability capabilities of the proposed methodology.

Overall, the results remain favourable across all three scenarios, with low MAE and prediction bounds that are not overly conservative. The figures also show each battery's state-space trajectory, where the circle marks the start of training, the square the start of online prediction, and the star the estimated EOL crossing. Averaging the nominal RUL and bound widths across batteries provides a population-level view of the mean prediction and uncertainty. To assess robustness against partial observability, two hard sensor failure scenarios are considered. At cycle $k_f = 55$, either the impedance sensor (x) or the SOH sensor (s) is assumed to fail. The missing state is then reconstructed online by forward integration of the nominal drift field, initialised from the last valid measurement prior to failure. Since the RUL value function is evaluated offline, the online query remains operational under both conditions. Figs. 7 and 6 show that the HJB framework maintains meaningful RUL estimates post-failure. For a fixed HJB value function, γ is varied to evaluate its effect on RUL predictions across all the fleet. Results as shown in Fig. 9 detail its sensitivity. The choice of γ has a linear effect on RUL prediction accuracy as shown in the figure. Predictions remain reliable only within range $\gamma \leq 0.2$, beyond which the fixed HJB value function cannot compensate for the shifted terminal condition. As error sensitivity $d(\text{MAE})/d\gamma$ consistently exceeds prediction sensitivity S_γ for $\gamma > 0.2$, this confirms that the boundary function must be treated as a calibrated model parameter rather than a free design choice, and that re-solving the HJB PDE per boundary configuration is necessary for robust RUL estimation.

5.1. Broader practical Implications.

Beyond the numerical evaluation, some practical implications are discussed. This preliminary work highlights a key sensitivity of the HJ reachability framework albeit with its illustrated prowess in terms of interpretability and robustness. RUL bounds may fail to capture batteries whose degradation trajectories deviate significantly from the fleet behavior used to construct the drift field. This is evidenced by B0018 and B0006, where atypical charge-transfer resistance trends and sparse state-space coverage in the high-impedance region lead to poorly supported drift estimates and, consequently, bounds that fail to encapsulate the true RUL. With a sufficiently large and diverse fleet, an anomaly detection or outlier rejection stage could be applied prior to drift field construction to exclude atypical trajectories and improve bound reliability.

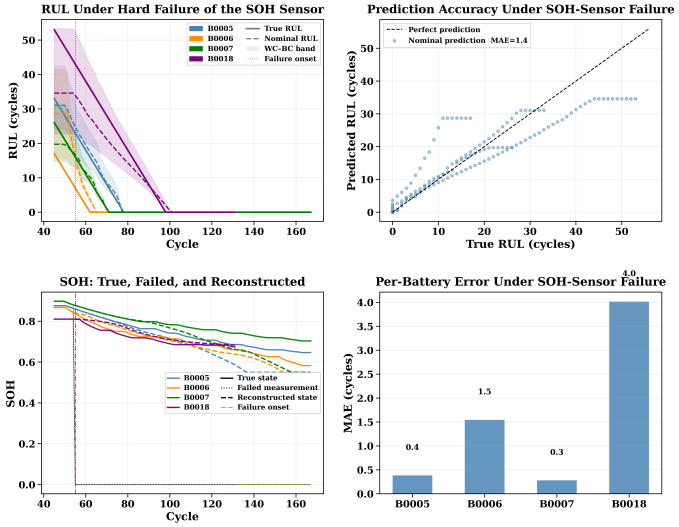

Figure 6. Robustness of the proposed scheme under SOH-input sensor loss. The panels show the RUL predictions after cycle 55 (failure onset), together with uncertainty bounds and true RUL; the predicted nominal RUL versus the true RUL; reconstruction based on the learned drift model using the last known reading; and the MAE after sensor failure onset.

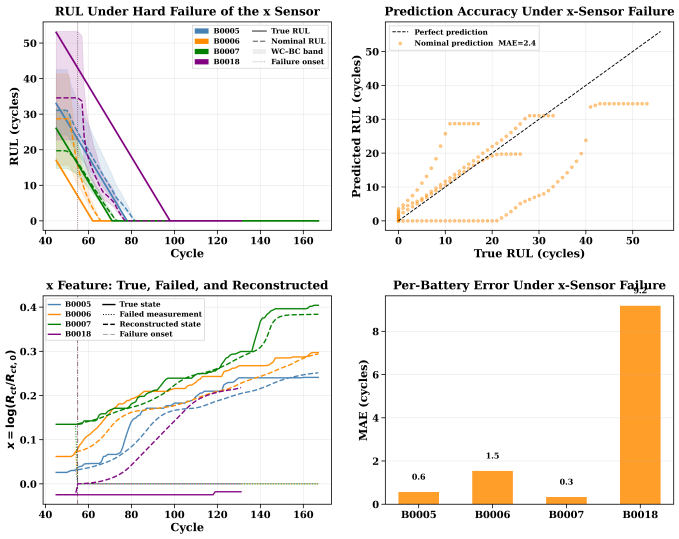

Figure 7. Robustness of the proposed scheme under impedance-input sensor loss. The panels show the RUL predictions after cycle 55 (failure onset), together with uncertainty bounds and true RUL; the predicted nominal RUL versus the true RUL; reconstruction based on the learned drift model using the last known reading; and the MAE after sensor failure onset.

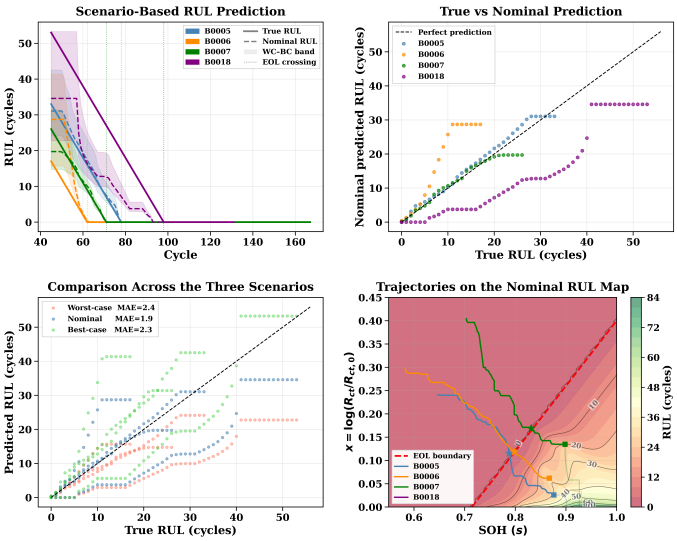

Figure 8. Evaluation of the proposed scheme. From top left to bottom right: the RUL predictions from cycle 46 onward, together with the uncertainty bounds and the true RUL; the predicted nominal RUL compared with the true RUL; the three prediction scenarios (best-case, nominal, and worst-case); and the degradation trajectories projected onto the RUL map.

6. CONCLUSION AND FUTURE WORK

This work presented an HJB-based framework for battery RUL prediction using a low-dimensional degradation state defined by SOH and impedance growth. The method provides both nominal RUL estimates and uncertainty bounds through learned drift fields, and the results show that it captures the main degradation behavior of the NASA cells while remaining interpretable through an explicit failure boundary in the state space. The sensor-failure analysis also shows that the framework remains usable under loss of one sensing channel, since the missing state can be propagated using the learned degradation dynamics. Performance remains strong under SOH sensor failure, while larger errors under impedance sensor failure highlight the need for improved uncertainty treatment during state reconstruction. Given that this is a preliminary work in the use of HJB in prognostics, authors agree that much work is required to build on this thesis. Thus future work will focus on these points: Incorporating measurement noise more explicitly, either directly in the state propagation or through a combination with filter-based state estimation methods such as Kalman or particle filtering. The use of richer and more diverse battery datasets will be explored to provide a stronger foundation for UQ in the RUL maps. A main concern in this work is a lack of consideration of the explicit operating conditions such as temperature and their effect on the variability of the assumed similar chemistry degradation trends. A future work will make such considerations. Finally, since the predicted RUL is strongly dependent on the chosen failure boundary, a more formal and principled

definition of this boundary will also be investigated in future work.

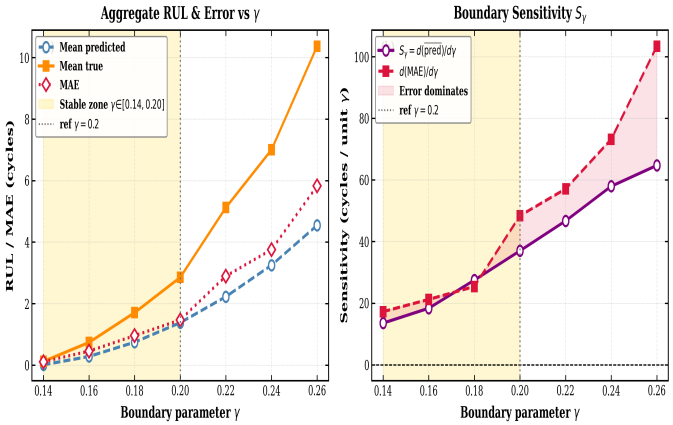

Figure 9. Sensitivity of HJB-predicted RUL to the EOL boundary parameter γ. (*Left*) Mean predicted RUL, true RUL, and MAE across GROUP6 batteries, beyond $\gamma = 0.2$ prediction error grows faster than predicted RUL. (*Right*) Error sensitivity $d(\text{MAE})/d\gamma$ consistently exceeds prediction sensitivity $S_\gamma = d(\overline{\text{pred}})/d\gamma$ for $\gamma > 0.2$.

ACKNOWLEDGMENT

This work is funded by Projet Electro'Mob, IMT Nord Europe, France.

REFERENCES

Al-Mohamad, A., Puig, V., & Hoblos, G. (2023). Recursive zonotopic set-membership approach for system-level prognostics with application to linear parameter-varying systems. *ISA Transactions*, *135*, 244–260.

Bansal, S., Chen, M., Herbert, S. L., & Tomlin, C. J. (2017). Hamilton-jacobi reachability: A brief overview and recent advances. In *2017 ieee 56th annual conference on decision and control (cdc)* (pp. 2242–2253).

Basora, L., Viens, A., Arias Chao, M., & Olive, X. (2025). A benchmark on uncertainty quantification for deep learning prognostics. *Reliability Engineering & System Safety*, *253*, 110513.

Boutrous, K., Puig, V., & Nejjari, F. (2022). A set-based prognostics approach for wind turbine blade health monitoring. *IFAC-PapersOnLine*, *55*(6), 402–407.

Boutrous, K., Thuillier, J., Jha, M. S., Puig, V., & Theilliol, D. (2023). Assessing a statistical and a set-based approach for remaining useful life prediction. In *2023 31st mediterranean conference on control and automation (med)* (pp. 25–30). doi: 10.1109/MED59994.2023.10185887

Ding, Z.-Q., Qin, Q., Zhang, Y.-F., & Lin, Y.-H. (2024). An uncertainty quantification and calibration framework for rul prediction and accuracy improvement. *IEEE Transactions on Instrumentation and Measurement*, *73*, 1–13. doi: 10.1109/TIM.2024.3485392

Helen, A., Huan, X., Paulson, N., et al. (2024). Probabilistic machine learning for battery health diagnostics and prognostics—review and perspectives. *npj Materials Sustainability*, *2*, 14. doi: 10.1038/s44296-024-00011-1

Jose, S., Zemouri, R., Khanh, N., Medjaher, K., Levesque, M., & Tahan, A. (2025, 05). Prognostics of complex machinery with sparse multilabel multimodal run-to-failure data: A graph neural network approach. *Advanced Engineering Informatics*, *65*, 103361. doi: 10.1016/j.aei.2025.103361

Jouin, M., Gouriveau, R., & Hissel, D. (2016). Particle filter-based prognostics: Review, discussion and perspectives. *Mechanical Systems and Signal Processing*, *72–73*, 2–31. Retrieved from http://dx.doi.org/10.1016/j.ymssp.2015.11.008 doi: 10.1016/j.ymssp.2015.11.008

Lin, Y.-H., Yan, P.-C., & Zio, E. (2025, 12). Recent advances in uncertainty analysis for prognostics and remaining useful life prediction: A review. *Reliability Engineering System Safety*, *269*, 112110. doi: 10.1016/j.ress.2025.112110

Nemani, V., Biggio, L., Huan, X., Hu, Z., Fink, O., Tran, A., ... Hu, C. (2023, December). Uncertainty quantification in machine learning for engineering design and health prognostics: A tutorial. *Mechanical Systems and Signal Processing*, *205*, 110796. doi: 10.1016/j.ymssp.2023.110796

Nunes, T. S. N., Moura, J. J. P., Prado, O. G., Camboim, M. M., Rosolem, M. d. F. N., Beck, R. F., ... Ding, H. (2023, June). An online unscented kalman filter remaining useful life prediction method applied to second-life lithium-ion batteries. *Electrical Engineering*, *105*, 3481–3492. doi: 10.1007/s00202-023-01910-7

Saha, B., & Goebel, K. (2007). *Battery data set.* NASA Ames Prognostics Data Repository. Moffett Field, CA. Retrieved from https://www.nasa.gov/intelligent-systems-division/discovery-and-systems-health/pcoe/pcoe-data-set-repository

APPENDIX

A. HAMILTON–JACOBI CHARACTERIZATION OF THE RUL VALUE FUNCTION

Let

$$V(z_0) := \inf\{t \geq 0 : z(t; z_0) \in \mathcal{T}_{\mathrm{EOL}}\}, \qquad \dot{z}(t) = f(z(t)).$$

Proposition 1. *If V is differentiable on $\Omega \setminus \mathcal{T}_{\mathrm{EOL}}$, then*

$$\nabla V(z) \cdot f(z) = -1, \qquad z \in \Omega \setminus \mathcal{T}_{\mathrm{EOL}},$$

with $V(z) = 0$ on $\mathcal{T}_{\mathrm{EOL}}$.

Proof. For sufficiently small $\Delta t > 0$,

$$V(z) = \Delta t + V(z(\Delta t)).$$

Thus

$$V(z(\Delta t)) - V(z) = -\Delta t.$$

Since

$$z(\Delta t) = z + f(z)\Delta t + o(\Delta t),$$

a first-order expansion gives

$$V(z(\Delta t)) - V(z) = \nabla V(z) \cdot f(z)\, \Delta t + o(\Delta t).$$

Dividing by Δt and letting $\Delta t \to 0$ yields

$$\nabla V(z) \cdot f(z) = -1.$$

Moreover, $V(z) = 0$ on $\mathcal{T}_{\mathrm{EOL}}$ by the definition of the boundary time. □