Journal of Intelligent Manufacturing and Special Equipment

207

Physics-informed digital twins for sustainability: a data-efficient framework for physically consistent state tracking

Aswin Karkadakattil
Department of Mechanical, Independent Researcher, Kanhangad, India

Received 29 December 2025
Revised 28 January 2026
22 February 2026
13 March 2026
Accepted 7 April 2026

Abstract

Purpose – This study presents a data-efficient and physically constrained digital twin framework for monitoring degradation and performance loss in engineered systems, with emphasis on sustainability-oriented applications where sensing availability, historical data, and computational resources may be limited. The objective is to establish a structurally interpretable and physically consistent approach to long-term degradation state tracking under sparse data conditions.

Design/methodology/approach – The proposed digital twin represents system sustainability using a compact internal degradation state governed by irreversible kinetic principles. Physical constraints, including monotonicity, boundedness, and asymptotic saturation, are embedded directly into the learning formulation to ensure physically admissible evolution. A lightweight neural network surrogate approximates the state trajectory while remaining constrained by the governing physical law. The framework is evaluated using five-fold cross-validation and validated against publicly available lithium-ion battery degradation datasets to assess robustness, reproducibility, and experimental consistency.

Findings – The results demonstrate stable and physically admissible state evolution across all validation folds, with degradation trajectories remaining monotonic, bounded, and insensitive to data partitioning. The physics-informed formulation achieves predictive accuracy comparable to purely data-driven approaches while preserving interpretability and structural consistency. Experimental validation confirms the ability of the framework to capture realistic degradation behaviour under sparse sampling without producing non-physical artifacts.

Originality/value – The framework embeds irreversible physical structure directly into a compact and identifiable degradation-state representation, prioritizing physical consistency, interpretability, and data efficiency. This provides a transparent and extensible foundation for sustainability-oriented digital twin implementation in intelligent engineering systems.

Funding: This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Competing interests: The author declares that there are no known financial or non-financial competing interests that could have influenced the work reported in this manuscript.

Journal of Intelligent Manufacturing and Special Equipment
Vol. 7 No. 2, 2026
pp. 207-239
Emerald Publishing Limited
e-ISSN: 2633-660X
p-ISSN: 2633-6596
DOI 10.1108/JIMSE-12-2025-0037

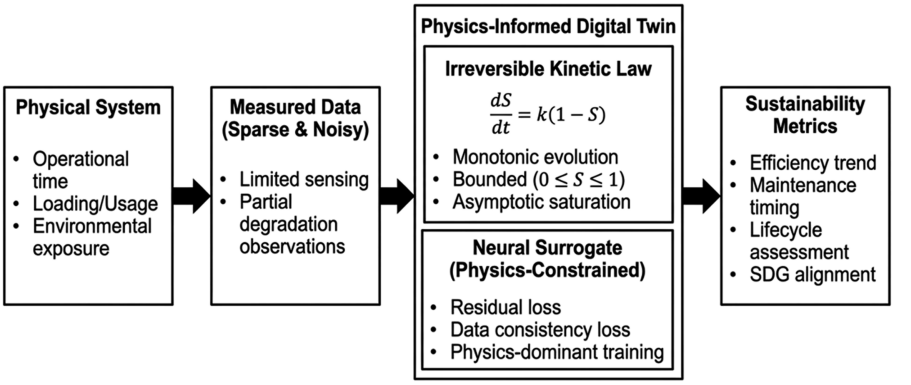

"Compact, interpretable, and physically admissible digital twin under data scarcity."

Graphical abstract

Keywords Digital twin, Physics-informed learning, Sustainability, Materials informatics, Smart manufacturing, Data-efficient modelling, Sustainable development goals
Paper type Research article

1. Introduction

Engineering systems that support modern infrastructure, manufacturing, transportation, and energy technologies are increasingly required to operate under stringent sustainability constraints. In addition to maintaining performance and reliability, these systems must minimize energy consumption, reduce material waste, and extend operational lifespan to align with global sustainability objectives (Isaksson *et al.*, 2015; Bertoni, 2017; Zetterlund *et al.*, 2016). Achieving these goals depends critically on the ability to monitor degradation, efficiency loss, and irreversible damage accumulation throughout the operational lifetime of engineering assets (Plota and Mašek, 2020; Arshad and Maaroufi, 2018). However, direct measurement of degradation or internal health states during operation is often impractical, costly, or technically infeasible, particularly for long-lived or safety-critical systems (Celina *et al.*, 2005; Roberts, 2003). Digital twins have emerged as a promising approach for addressing this challenge by providing virtual representations of physical systems that evolve alongside their real counterparts (Madni *et al.*, 2019; Wagg *et al.*, 2020; Bado *et al.*, 2022). These digital representations enable continuous monitoring, predictive maintenance, and lifecycle-aware decision-making across diverse engineering applications (Iranshahi *et al.*, 2025; Ren *et al.*, 2025). When implemented effectively, digital twins can improve operational efficiency, enhance reliability, and support sustainability-oriented asset management. However, many existing digital twin implementations rely on high-fidelity numerical simulations or large volumes of experimental and operational data (Cuomo *et al.*, 2022; Lawal *et al.*, 2022). While such approaches can provide detailed system representation, they are often computationally expensive, data-intensive, and difficult to deploy in resource-constrained or data-limited environments (Pannala *et al.*, 2024; Shin *et al.*, 2024). These limitations are particularly relevant in sustainability-oriented monitoring applications, where scalable, interpretable, and data-efficient degradation tracking methods are required. In parallel, degradation modelling has been extensively studied using structured state-space models and stochastic process formulations. Classical approaches include Wiener processes, Gamma processes, inverse Gaussian processes, and related cumulative damage models (Abdel-Hameed, 1975; van Noortwijk, 2009; Gorjian *et al.*, 2010; Ye and Chen, 2014; Ye *et al.*, 2015).

These models provide flexible probabilistic descriptions of degradation evolution and uncertainty, and have been widely applied in reliability engineering, maintenance planning, and prognostics. In particular, Gamma process formulations inherently ensure non-negative and monotonic degradation increments, while constrained Wiener process variants can represent irreversible degradation through positive drift structures. Such stochastic models have proven effective for statistical lifetime prediction and uncertainty quantification. However, these approaches are primarily probabilistic in nature and typically rely on empirical parameterization of degradation increments or drift terms. They do not explicitly derive degradation evolution from mechanistic irreversible kinetic principles, nor are they generally embedded within physics-informed learning architectures that structurally enforce thermodynamic admissibility. As a result, while stochastic models can capture observed degradation trends and ensure monotonicity under appropriate formulations, their integration into digital twin frameworks for sustainability-oriented monitoring does not inherently guarantee structural physical interpretability or analytically constrained state evolution. Recent advances in physics-informed learning have expanded the capabilities of digital twin frameworks by embedding governing physical relationships directly into machine learning models. Hybrid physics–data approaches have demonstrated improved robustness and interpretability in applications such as wind turbine degradation monitoring and industrial thermal system modelling (Yucesan and Felipe, 2023; Majumdar *et al*., 2025; Yang *et al*., 2024). These approaches represent significant progress in combining physical knowledge with data-driven learning. However, much of the existing work focuses on reconstructing high-dimensional physical fields or modelling system-specific mechanisms in detail. While such formulations can achieve high predictive fidelity, they often require extensive data, system-specific calibration, and substantial computational resources. Purely data-driven digital twin models have also been explored (Seyyedi *et al*., 2023; Pawar *et al*., 2021). Although these approaches can achieve accurate interpolation within the training domain, their lack of explicit physical constraints may lead to non-physical or unstable behaviour when extrapolated beyond observed data (Perdikaris, 2024; Anitescu *et al*., 2023). This limitation reduces reliability in applications where physically consistent degradation tracking over long time horizons is essential. Physics-informed learning provides a structured alternative by constraining model evolution through physically meaningful governing relationships (Pang *et al*., 2019; Klamecki, 1984; Estevez and Vilanova, 2009). By embedding irreversible physical principles directly into the model formulation, such approaches can ensure monotonic, bounded, and physically admissible degradation evolution while reducing dependence on dense datasets (Roberts, 2003; Celina *et al*., 2005). Nevertheless, relatively few digital twin formulations have focused specifically on sustainability-oriented degradation tracking using compact, structurally identifiable internal state representations governed by analytically constrained irreversible kinetics. Many existing approaches prioritize architectural complexity or system-specific modelling rather than structural physical consistency, analytical tractability, and interpretability. The primary novelty of this work lies in establishing a structurally constrained physics-informed digital twin framework in which degradation evolution is governed explicitly by physically admissible irreversible kinetics. Unlike conventional physics-informed neural network implementations that focus primarily on high-dimensional system reconstruction or predictive accuracy, the proposed formulation emphasizes structural physical admissibility, analytical transparency, and parameter identifiability. By representing degradation as an explicitly governed internal state rather than a purely empirical or latent stochastic variable, the framework ensures monotonic, bounded, and stable evolution consistent with irreversible physical processes. This approach provides a minimal yet physically grounded foundation for sustainability-oriented digital twin implementation, enabling reliable degradation tracking under sparse data conditions while preserving interpretability and computational efficiency. To address this gap, the present study introduces a physics-informed digital twin framework in which degradation and sustainability-related performance loss are represented through a compact internal state

governed by irreversible kinetics. The proposed formulation ensures physically admissible state evolution, including monotonic progression, boundedness, and asymptotic stability, while remaining computationally efficient and interpretable. By embedding physically grounded degradation dynamics directly into the learning formulation, the framework provides a scalable and structurally consistent approach for sustainability-oriented digital twin implementation. This formulation is particularly suitable for applications such as intelligent manufacturing systems, equipment degradation monitoring, battery aging assessment, and lifecycle-aware asset management, where reliable and physically interpretable degradation tracking is essential for sustainable operation and decision support.

2. Conceptual framework of the physics-informed digital twin

The physics-informed digital twin developed in this study is designed to enable reliable and interpretable monitoring of degradation and performance loss under limited data conditions. Rather than attempting to reproduce the full Multiphysics complexity of an engineering system, the proposed framework focuses on capturing the dominant irreversible processes that govern long-term degradation, efficiency loss, and entropy generation (Klamecki, 1984; Celina *et al*., 2005). This design choice is intentional. Many existing digital twin implementations prioritize detailed physical modelling or computational sophistication. While such approaches can provide highly accurate system representations, they often require extensive parameterization, large datasets, and significant computational resources. In contrast, the present framework emphasizes structural physical consistency, interpretability, and scalability. At the core of the framework is a compact internal state variable that represents cumulative degradation or sustainability-related performance loss. This internal state evolves according to physically motivated irreversible kinetics, ensuring monotonic progression and bounded behavior consistent with thermodynamic principles (Roberts, 2003; Estevez and Vilanova, 2009). By embedding these physical constraints directly into the model formulation, physically admissible system evolution is guaranteed by design rather than learned implicitly from data. The digital twin receives operational observations, such as time history, usage patterns, or environmental exposure conditions. These observations serve as informative signals that guide the digital twin while remaining consistent with the governing physical constraints. This structure improves robustness under sparse, noisy, or incomplete data conditions, which are common in real-world engineering applications (Lawal *et al*., 2022; Pannala *et al*., 2024). A neural network component is used to approximate the evolution of the internal degradation state while remaining constrained by the governing physical law. Importantly, the neural network does not replace the physics, but rather adapts the physically constrained state trajectory to available observations. This hybrid approach combines the flexibility of machine learning with the stability and interpretability of physically governed evolution. The resulting digital twin provides a transparent and computationally efficient representation of system degradation that can support sustainability-oriented decision-making. The internal degradation state can be directly interpreted in terms of system health, efficiency trends, maintenance requirements, and lifecycle progression. This enables reliable long-term monitoring while avoiding the complexity and data requirements associated with high-fidelity simulation-based digital twins. The overall framework integrates physical constraints, data-driven learning, and sustainability-oriented state interpretation into a unified structure. By prioritizing physical consistency, interpretability, and data efficiency, the proposed approach provides a scalable and extensible foundation for digital twin development in intelligent engineering systems.

2.1 Objectives and contributions of the present work

The objectives and contributions of the present study are summarized as follows:

(1) *Physically Consistent Degradation-State Representation* To develop a digital twin formulation in which degradation and sustainability-related performance loss are represented through a compact internal state governed by irreversible physical principles, ensuring monotonic and bounded evolution.

(2) *Analytical Consistency and Structural Guarantees* To establish and verify the fundamental analytical properties of the degradation state, including stability, boundedness, and structural identifiability, within a physics-informed learning framework.

(3) *Data-Efficient Learning Under Sparse Observations* To demonstrate that physically constrained degradation dynamics can be learned reliably using limited and sparsely sampled data, without requiring high-fidelity simulation models or extensive training datasets.

(4) *Scalable Framework for Sustainable Engineering Applications* To provide a computationally efficient and interpretable digital twin framework suitable for intelligent manufacturing systems and sustainability-oriented engineering applications, with clear pathways for domain-specific calibration and deployment.

As shown in Figure 1, the proposed physics-informed digital twin framework captures the interaction between the physical system and its digital counterpart, enabling sustainability-aware decision-making.

3. Physics-informed digital twin formulation

3.1 Scalability and generalization considerations

The proposed digital twin framework is initially formulated using a single internal degradation state to provide a clear, physically interpretable, and structurally identifiable representation of system evolution. This low-dimensional formulation is an intentional design choice aimed at establishing a minimal yet physically consistent baseline that remains robust under limited data conditions. By representing degradation through a compact internal variable, the

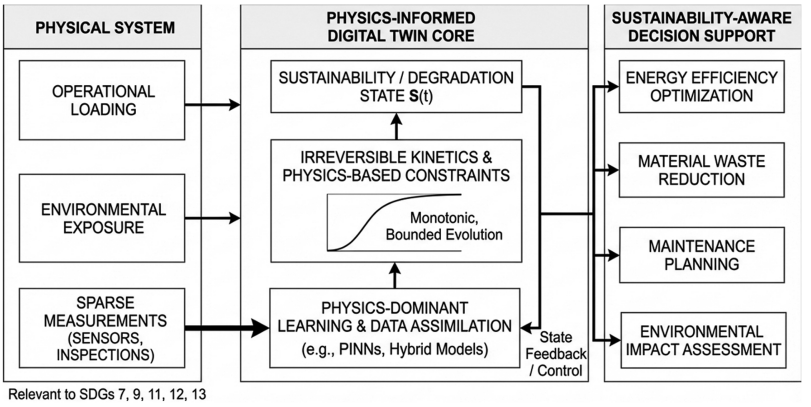

Figure 1. Conceptual architecture of the physics-informed digital twin for sustainability. Schematic representation of the proposed physics-informed digital twin framework, showing the interaction between the physical system and the digital twin, the physics-governed evolution of an internal sustainability state, and the translation of this state into sustainability-aware decision support. The framework emphasizes physical consistency, data efficiency, and interpretability rather than high-fidelity numerical simulation

framework reduces parameter redundancy and improves stability during training, particularly when observations are sparse, noisy, or unevenly distributed. The single-state representation serves as a foundational model that captures the essential characteristics of irreversible degradation, including monotonic progression, bounded evolution, and asymptotic stability. These structural properties arise directly from the imposed physical constraints and ensure that the learned state remains physically meaningful throughout the system lifecycle. Importantly, this formulation provides analytical transparency, allowing the relationship between physical assumptions and model behavior to be clearly understood. Although the baseline model uses a single degradation state, the framework is inherently modular and extensible. Many real engineering systems exhibit multiple interacting degradation mechanisms operating across different temporal or spatial scales. To represent such complexity, the formulation can be extended by introducing additional physically governed internal states, each corresponding to a distinct degradation process or subsystem. These extensions enable systematic expansion to multi-state digital twin architectures while preserving the underlying physical constraints that ensure irreversible and bounded evolution. This modular structure allows the framework to be adapted to application-specific requirements without altering its core physical principles. Higher-dimensional state representations, spatially distributed models, or multi-fidelity formulations can be incorporated as needed while maintaining interpretability and physical consistency. In this sense, the single-state implementation should be viewed not as a limitation, but as a structurally transparent starting point that establishes the fundamental physical and analytical properties of the proposed digital twin. More complex models can be developed as extensions of this baseline when additional system detail or representational capacity is required.

3.2 Physics-informed sustainability state formulation

System sustainability and degradation are represented through a compact internal state variable, denoted as $\tau(t)$, which evolves continuously over operational time. This internal state captures the cumulative effects of degradation, efficiency loss, and irreversible dissipation occurring during system operation. Rather than explicitly resolving individual degradation mechanisms, which may be complex or difficult to observe directly, $\tau(t)$ provides a physically interpretable macroscopic representation of long-term system evolution suitable for monitoring, analysis, and sustainability-oriented decision support.

The evolution of the sustainability state is governed by an irreversible kinetic relationship of the form:

$$\frac{d\tau}{dt} = k(1-\tau), k > 0$$

where k is a degradation rate constant that characterizes the characteristic timescale of irreversible system evolution.

This formulation is not arbitrary; it represents the minimal physically admissible evolution law satisfying the fundamental thermodynamic and stability requirements of irreversible degradation processes. Specifically, several essential physical properties are guaranteed by construction:

3.2.1 Irreversibility. Since $k > 0$ and $(1-\tau) \geq 0$ for $\tau \in [0, 1]$,

$$\frac{d\tau}{dt} \geq 0$$

This ensures that the degradation state evolves monotonically and does not decrease over time, consistent with the irreversible nature of entropy production and cumulative damage processes in physical systems.

3.2.2 Boundedness. The sustainability state is defined within the normalized interval

$$\tau(t) \in [0, 1]$$

where $\tau = 0$ represents the undegraded condition and $\tau = 1$ represents the asymptotic degradation limit. This bounded representation prevents non-physical divergence and ensures stable long-term behavior.

3.2.3 Asymptotic saturation. As degradation progresses and $\tau \to 1$, the evolution rate decreases:

$$\lim_{\tau \to 1} \frac{d\tau}{dt} = 0$$

This behavior reflects the gradual slowing of degradation commonly observed in aging processes, where systems approach a degraded equilibrium state.

From a thermodynamic perspective, this formulation is consistent with the principle of non-negative entropy production in irreversible systems. The monotonic increase of τ(t) represents cumulative dissipative effects, while boundedness ensures physically realistic long-term evolution. These structural properties arise directly from the governing kinetic relationship rather than being inferred implicitly from data. An additional advantage of this formulation is structural identifiability. The degradation dynamics are governed by a single physically meaningful parameter *k*, which reduces parameter ambiguity and improves interpretability. Unlike high-dimensional or purely data-driven models, where multiple parameter combinations may produce similar outputs, the present formulation ensures a unique and physically interpretable representation of degradation behavior. Within the physics-informed digital twin framework, the neural network serves as a constrained function approximator that learns the evolution of τ(t) while remaining consistent with the governing physical law. Importantly, the neural network does not alter the underlying kinetic structure but instead enables flexible adaptation to observed data while preserving physical admissibility. The first-order irreversible kinetic formulation therefore provides a physically grounded and analytically transparent baseline for sustainability state evolution. While more complex degradation models may be required for specific applications, this minimal formulation establishes a structurally consistent foundation that ensures stability, interpretability, and thermodynamic admissibility. System-specific calibration can be performed through estimation of the degradation rate parameter using operational or experimental data, and the formulation can be extended when additional complexity is required.

3.2.4 Physical and thermodynamic justification of the degradation kinetics. The irreversible kinetic formulation adopted in this study is consistent with established physical and thermodynamic principles governing degradation processes in engineered systems. Degradation in materials and functional components arises from irreversible microstructural and physicochemical changes, including defect accumulation, fatigue damage, corrosion, and electrochemical side reactions. These processes progressively reduce system performance and cannot spontaneously reverse under normal operating conditions. As a result, degradation evolution must be represented using a physically admissible internal state variable that evolves monotonically with time.

Let $\tau(t) \in [0, 1]$ denote a normalized internal degradation state, where $\tau = 0$ corresponds to the undegraded condition and $\tau = 1$ represents the asymptotic degradation limit. A physically admissible degradation model must satisfy the irreversibility condition:

$$\frac{d\tau}{dt} \geq 0$$

ensuring that cumulative degradation does not decrease over time.

More generally, degradation evolution can be expressed in the form:

$$\frac{d\tau}{dt} = f(\tau)$$

where the function $f(\tau)$ must satisfy the conditions:

$$f(\tau) \geq 0 \text{for } \tau \in [0, 1], f(1) = 0$$

These conditions ensure both irreversibility and bounded evolution of the degradation state.

The simplest continuously differentiable function satisfying these physical admissibility requirements is the first-order kinetic relationship:

$$\frac{d\tau}{dt} = k(1 - \tau), k > 0$$

This formulation ensures that degradation progresses monotonically while approaching a stable limiting value over time. Physically, this behavior reflects the gradual reduction in degradation rate as the system approaches its asymptotic degraded state, a phenomenon commonly observed in fatigue, material aging, and electrochemical degradation processes. The first-order kinetic formulation is consistent with established degradation modelling approaches, including continuum damage mechanics evolution laws, empirical aging models, and reaction-kinetics-based degradation descriptions. Similar mathematical structures have been used extensively in degradation modelling literature, including damage evolution models in materials (Lemaitre, 2012) and electrochemical degradation studies (Celina *et al*., 2005). In the present framework, this kinetic relationship provides a physically interpretable and mathematically well-posed representation of degradation evolution. Importantly, it represents the minimal structurally admissible formulation capable of ensuring irreversible, bounded, and stable degradation state evolution while remaining suitable for integration within a physics-informed digital twin architecture.

3.2.5 Analytical solution and stability analysis. The degradation evolution equation introduced in Section 3.2,

$$\frac{d\tau}{dt} = k(1 - \tau), k > 0$$

is a first-order differential equation that admits a closed-form analytical solution. For an initial degradation state $\tau(0) = \tau_0$, the solution can be expressed as:

$$\tau(t) = 1 - (1 - \tau_0)e^{-kt}$$

This analytical expression provides clear insight into the long-term behavior and physical consistency of the degradation state.

First, the solution evolves monotonically with time. Because the degradation rate constant k is positive and the term $(1 - \tau)$ remains non-negative within the admissible range $\tau \in [0, 1]$, the rate of change $\frac{d\tau}{dt}$ remains non-negative. This ensures that the degradation state cannot decrease over time, reflecting the irreversible nature of physical degradation processes.

Second, the degradation state remains bounded within physically meaningful limits. For any admissible initial condition $\tau_0 \in [0, 1]$, the analytical solution satisfies:

$$0 \leq \tau(t) \leq 1 \text{for all } t \geq 0$$

This bounded behavior prevents divergence or non-physical values and ensures stable long-term evolution.

Third, the degradation state approaches a stable limiting value as time increases. Specifically,

$$\lim_{t\to\infty} \tau(t) = 1$$

This asymptotic convergence represents gradual saturation of degradation, a behavior commonly observed in aging, fatigue, and electrochemical degradation processes. As the system approaches this limit, the degradation rate decreases smoothly, ensuring stable and physically realistic behavior. In addition to these physical properties, the governing differential equation is mathematically well-posed. The evolution law is continuous and satisfies standard existence and uniqueness conditions, ensuring that the degradation trajectory is uniquely determined by its initial state. Overall, the existence of a closed-form analytical solution confirms that the proposed degradation formulation provides a stable, bounded, and physically consistent representation of irreversible system evolution. These analytical properties also support reliable integration of the degradation model within the physics-informed digital twin framework.

3.3 Extension beyond first-order kinetics for multi-phase degradation

While the first-order irreversible kinetic model provides a physically interpretable and analytically transparent baseline, many engineering systems exhibit more complex degradation behavior. These may include gradual transitions between degradation regimes, interaction-driven damage accumulation, or accelerated wear-out near the end of the system lifecycle. Such behavior is commonly observed in fatigue processes, material aging, and electrochemical degradation, where degradation rates evolve over time rather than remaining constant. To represent these effects while preserving physical admissibility, the governing evolution law can be extended to include nonlinear interaction terms. A generalized degradation model can be expressed as:

$$\frac{d\tau}{dt} = k_1(1-\tau) + k_2\tau(1-\tau),$$

where $k_1 > 0$ represents the baseline irreversible degradation rate and $k_2 \geq 0$ captures interaction-dependent or state-dependent degradation effects. The additional nonlinear term allows the degradation rate to vary as a function of the current degradation state, enabling representation of gradual changes in degradation dynamics while preserving the fundamental physical constraints.

This generalized formulation retains the key structural properties required for physically admissible degradation modelling:

3.3.1 Monotonicity. For $\tau \in [0, 1]$ and $k_1, k_2 \geq 0$,

$$\frac{d\tau}{dt} \geq 0$$

ensuring irreversible progression of degradation without artificial recovery.

3.3.2 Boundedness. The degradation state remains confined to the physically meaningful interval:

$$\tau(t) \in [0, 1]$$

preventing non-physical divergence or instability.

3.3.3 Asymptotic stability. As $\tau \to 1$, the degradation rate approaches zero:

$$\lim_{\tau \to 1} \frac{d\tau}{dt} = 0$$

ensuring smooth and stable convergence toward the degradation limit.

The nonlinear interaction term provides additional flexibility in representing state-dependent degradation dynamics while preserving thermodynamic admissibility and structural stability. For systems exhibiting clearly distinguishable lifecycle phases, such as early-stage activation, stable operation, and accelerated wear-out, a multi-state formulation can be introduced. In this approach, degradation is represented through multiple physically governed internal variables corresponding to different degradation processes:

$$\frac{d\tau_1}{dt} = k_a(1 - \tau_1),$$
$$\frac{d\tau_2}{dt} = k_b\tau_1(1 - \tau_2),$$

where τ_1 represents early-stage activation or conditioning effects, and τ_2 represents cumulative wear progression. This formulation enables representation of sequential or interacting degradation mechanisms commonly observed in engineering systems.

The overall degradation state is then defined as a weighted combination of individual degradation components:

$$\tau = w_1\tau_1 + w_2\tau_2,$$

where $w_1 \geq 0$, $w_2 \geq 0$, and $w_1 + w_2 = 1$. This weighted representation preserves boundedness and provides a physically interpretable measure of cumulative system degradation. These extended formulations preserve the core structural principles of irreversible, bounded, and thermodynamically consistent degradation evolution while enabling representation of more complex degradation patterns. Importantly, the extensions remain fully compatible with the physics-informed learning framework and can be implemented using the same constrained learning structure as the baseline model. The extended models should be interpreted as physically admissible generalizations that increase representational flexibility while maintaining interpretability, stability, and structural consistency. The appropriate model complexity can be selected based on application requirements, data availability, and system characteristics.

3.3.4 Mathematical and physical interpretation of the nonlinear degradation extension. While the first-order irreversible kinetic formulation provides a physically admissible and analytically transparent baseline, many engineering systems exhibit degradation behavior in which the degradation rate evolves with the current state of accumulated damage. Such behavior may arise from interacting degradation mechanisms, microstructural evolution, material weakening, or progressive loss of structural integrity. To represent these effects while preserving physical consistency and mathematical admissibility, the degradation evolution equation can be extended to include a state-dependent term:

$$\frac{d\tau}{dt} = k_1(1 - \tau) + k_2\tau(1 - \tau)$$

where $k_1 > 0$ represents the baseline degradation rate and $k_2 \geq 0$ introduces state-dependent modulation of the degradation kinetics.

This formulation can be rewritten as:

$$\frac{d\tau}{dt} = (k_1 + k_2\tau)(1 - \tau)$$

which provides a clearer physical interpretation. The degradation rate remains proportional to the remaining undegraded fraction $(1-\tau)$, while the effective degradation rate coefficient $(k_1 + k_2\tau)$ varies smoothly with the current degradation state. This structure reflects the possibility that degradation kinetics may evolve as damage accumulates, consistent with experimental observations in fatigue damage accumulation, wear processes, and electrochemical aging, where degradation behavior depends on the current system condition.

Importantly, this extended formulation preserves the key physical admissibility requirements. For all admissible states $\tau \in [0, 1]$ and parameters $k_1 > 0$, $k_2 \geq 0$,

$$\frac{d\tau}{dt} \geq 0$$

ensuring irreversible degradation progression. Furthermore,

$$\frac{d\tau}{dt} = 0 \text{when} \tau = 1$$

which guarantees bounded evolution and prevents divergence beyond the physically meaningful degradation limit.

Depending on parameter values, the extended formulation may produce degradation trajectories with varying curvature, including concave or sigmoid-like profiles. Such behavior reflects state-dependent degradation kinetics and remains fully consistent with irreversible physical processes. The degradation state remains monotonic, bounded, and asymptotically stable, and no artificial recovery, oscillation, or instability can occur under admissible parameter conditions. It is important to note that this nonlinear formulation is presented as a general physically admissible extension rather than a universal degradation law. Its applicability depends on system-specific degradation mechanisms, and parameter values should be estimated using domain-relevant experimental or operational data. The inclusion of the state-dependent term therefore provides additional modelling flexibility while preserving the mathematical well-posedness, physical admissibility, and interpretability of the degradation evolution equation. Within the physics-informed digital twin framework, this extension enables representation of more complex degradation dynamics when required by specific applications, while maintaining the same physically grounded and structurally consistent formulation established by the baseline irreversible kinetic model. As shown in Figure 2, different degradation trajectories are compared across classical first-order kinetics, nonlinear saturation behaviour, and a two-state lifecycle formulation.

4. Implementation and validation methodology

4.1 Neural digital twin architecture

The physics-informed digital twin is implemented using a feedforward neural network that approximates the sustainability state $\tau(t)$ as a continuous and differentiable function of operational time. The neural network serves as a constrained function approximator, enabling flexible representation of degradation evolution while remaining consistent with the governing physical law. The architecture consists of fully connected layers with smooth nonlinear activation functions to ensure differentiability, which is required for enforcing the physics-based evolution equation. In the present implementation, the network includes two hidden layers with 32 neurons per layer and hyperbolic tangent (tanh) activation functions. The tanh activation was selected due to its smooth and continuously differentiable nature, which improves gradient stability and enables accurate enforcement of the governing differential

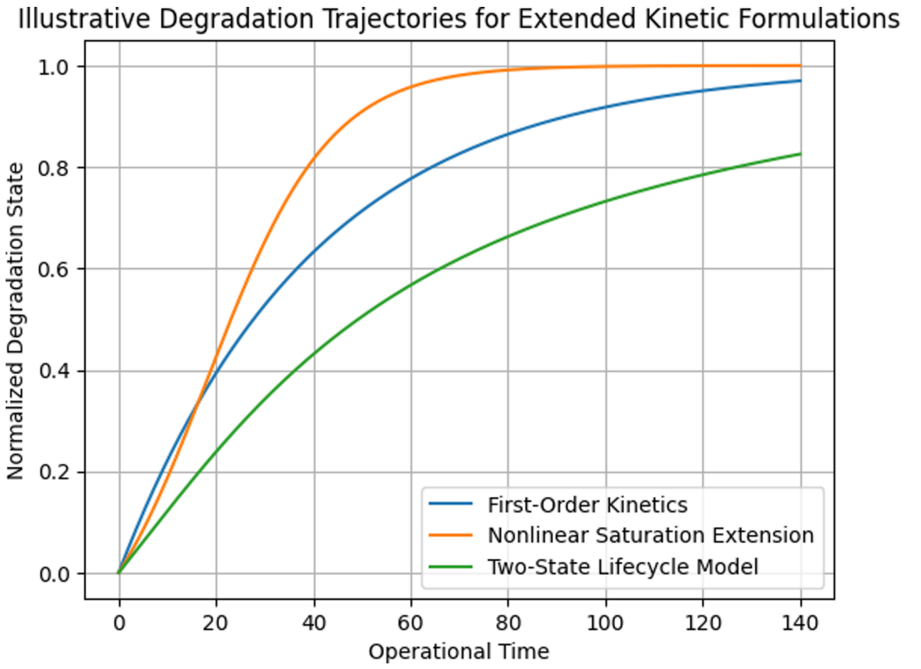

Figure 2. Illustrative degradation trajectories for extended kinetic formulations. Comparison between classical first-order irreversible kinetics, a nonlinear saturation model incorporating wear-out acceleration, and a two-state lifecycle formulation. All models preserve monotonic and bounded evolution of the normalized degradation state, demonstrating the structural extensibility of the proposed physics-informed digital twin framework

equation through automatic differentiation. This architecture was selected as a minimal configuration sufficient to ensure stable convergence while avoiding over-parameterization, consistent with the data-efficient design objective of the proposed framework. To ensure physical admissibility, a bounded sigmoid activation function is applied at the output layer, explicitly enforcing the constraint:

$$0 \leq \tau(t) \leq 1$$

This guarantees that predicted degradation states remain within physically meaningful limits during both training and inference. Importantly, the neural network does not replace the governing physical model. Instead, it operates within the constraints imposed by the irreversible kinetic law. The role of the neural network is to learn a smooth state trajectory that satisfies both the physical evolution equation and available observational data. This physics-constrained architecture improves stability, interpretability, and extrapolation reliability compared to purely data-driven neural networks, which may produce non-physical predictions when trained on limited datasets. Training was performed using the Adam optimization algorithm with a learning rate of 0.001. The network was trained for 2,500 epochs using full-batch gradient descent. These training settings provided stable convergence, consistent enforcement of physical constraints, and reproducible results across all validation folds.

4.2 Physics-informed loss function

Physical consistency of the digital twin is enforced through a composite loss function that integrates both physics-based constraints and observational data. This formulation ensures that

the learned degradation trajectory remains consistent with the governing irreversible kinetic law while maintaining agreement with experimentally observed degradation behaviour.

The sustainability state $\tau(t)$ is constrained to satisfy the irreversible kinetic evolution equation:

$$\frac{d\tau}{dt} = k(1-\tau),$$

where $k > 0$ represents the degradation rate constant.

To enforce adherence to this governing equation during training, a physics-residual loss term is defined as:

$$L_{\text{phys}} = \frac{1}{N_p}\sum_{i=1}^{N_p}\left(\frac{d\tau(t_i)}{dt} - k[1-\tau(t_i)]\right)^2,$$

where t_i denotes collocation points sampled across the operational time domain. This term penalizes deviations from the prescribed irreversible kinetic relationship and promotes physically admissible state evolution throughout the domain.

To incorporate available observational information, a data-consistency loss term is introduced:

$$L_{\text{data}} = \frac{1}{N_d}\sum_{j=1}^{N_d}\left(\tau(t_j) - \tau^{\text{obs}}(t_j)\right)^2,$$

where $\tau^{\text{obs}}(t_j)$ represents experimentally measured or reference degradation values.

The total training loss is defined as:

$$L = L_{\text{phys}} + \lambda L_{\text{data}},$$

where the parameter λ controls the relative weighting between physics constraint enforcement and empirical data fitting. Smaller values of λ yield behaviour closer to purely data-driven learning, while larger values place stronger emphasis on adherence to the governing irreversible kinetic constraint. In the present study, the value $\lambda = 0.1$ was selected empirically to achieve a balanced training process in which the governing physical constraint remains the dominant influence while still allowing adaptation to observed degradation data. This physics-prioritized training strategy promotes stable convergence, mitigates overfitting to sparse measurements, and ensures that the learned degradation trajectories remain monotonic, bounded, and physically interpretable across all validation folds.

4.2.1 Sensitivity analysis of the physics–data weighting parameter (λ). To examine the influence of the physics–data weighting parameter λ, a parametric sensitivity analysis was performed. As defined in Section 4.2, λ determines the balance between the physics-residual loss L_{phys}, which enforces the irreversible kinetic constraint, and the data-consistency loss L_{data}, which promotes agreement with observed degradation measurements. Consequently, λ directly affects the relative dominance of structural physical consistency versus empirical fitting during optimization.

The model was trained using

$$\lambda \in \{0.01, 0.05, 0.1, 0.5, 1.0\},$$

while maintaining identical network architecture, initialization, and training settings across all experiments.

Figure 3(a) presents the resulting degradation trajectories $\tau(t)$ for each λ value. Across the entire tested range, the learned trajectories remain monotonic, bounded, and asymptotically stable, reflecting structural enforcement of the irreversible kinetic evolution law. When λ is small (e.g. 0.01), training becomes more data-driven, resulting in slightly increased sensitivity during early-stage degradation. Conversely, larger λ values (e.g. 0.5 and 1.0) impose stronger physics regularization, producing smoother early-time evolution. Importantly, despite these

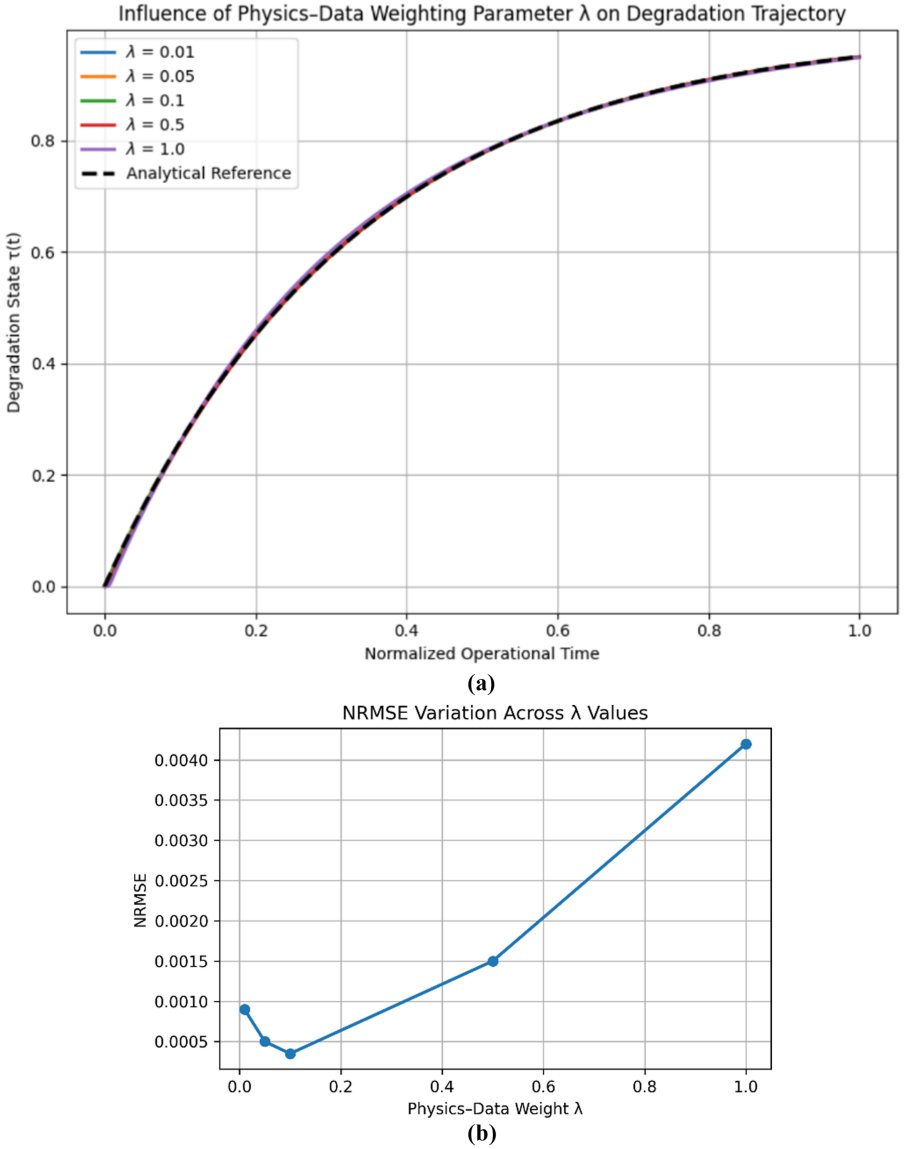

Figure 3. Sensitivity analysis of the physics–data weighting parameter λ. (a) Learned degradation trajectories for different λ values, showing preserved monotonic and bounded behaviour. (b) Variation of NRMSE with λ, indicating a balanced regime at $\lambda = 0.1$

variations, the overall degradation trends remain consistent due to the dominant role of the embedded physical formulation. Quantitative evaluation using the normalized root-mean-square error (NRMSE) is shown in Figure 3(b). The error exhibits a shallow U-shaped dependence on λ, with minimum NRMSE occurring at $\lambda = 0.1$, indicating an optimal trade-off between physics enforcement and data fidelity. For very small λ, reduced physics regularization leads to modestly increased prediction error. When λ becomes excessively large, over-regularization slightly limits data alignment, again increasing error. In addition to trajectory variation, convergence behaviour was assessed by monitoring loss stabilization during training. Larger λ values resulted in faster decay of the physics-residual component, whereas very small λ values required additional epochs to achieve consistent constraint satisfaction. No instability or oscillatory training behaviour was observed across the tested range, confirming structural robustness of the physics-informed formulation. Overall, the results demonstrate that the framework is robust to reasonable variations in λ, while achieving optimal predictive performance near $\lambda = 0.1$. This supports the selection of $\lambda = 0.1$ for all subsequent experiments.

4.3 Cross-validation strategy

Model robustness and generalization performance are evaluated using five-fold cross-validation. The dataset is divided into five equally sized subsets. For each fold, four subsets are used for training and one subset is reserved for validation. This process is repeated until each subset has served as the validation set. This validation strategy assesses the stability and consistency of the learned degradation state across different data partitions. Unlike conventional machine learning applications, where predictive accuracy is the primary objective, the focus here is on evaluating whether the learned state evolution remains physically consistent and structurally stable across all folds. Close agreement between degradation trajectories obtained from different folds indicates that the learned digital twin behavior is governed primarily by embedded physical constraints rather than sensitivity to specific data partitions.

4.4 Implementation transparency and reproducibility

To ensure full reproducibility, the proposed physics-informed digital twin was implemented using the PyTorch deep learning framework. The neural network architecture, training configuration, and optimization parameters were explicitly defined to ensure consistent and stable learning of the degradation state while preserving the embedded physical constraints. The model was trained using standard CPU hardware, demonstrating that the framework can be deployed without specialized computational resources. The neural network consists of fully connected layers with smooth activation functions to enable differentiable state evolution and enforcement of the governing differential equation. A bounded output activation function was used to ensure that the predicted degradation state remains within physically admissible limits. The model was trained using a physics-informed loss formulation combining the governing equation residual and data consistency terms, as described in Section 4.2. Training was performed across five cross-validation folds, and stable convergence was observed in all cases. The complete neural network architecture and training parameters used in this study are summarized in Table 1. These settings were selected to provide stable optimization while maintaining computational efficiency and reproducibility.

All implementation scripts, including model definition, training procedures, and figure generation routines, are structured to allow straightforward reproduction of the results. The modular design enables adaptation of the framework to alternative degradation datasets and engineering applications with minimal modification.

Table 1. Neural network architecture and training parameters

Parameter	Value
Framework	PyTorch
Network type	Fully connected neural network
Hidden layers	2
Neurons per hidden layer	32
Activation function	Tanh
Output activation	Sigmoid
Optimizer	Adam
Learning rate	0.001
Training epochs	2,500
Loss weight (λ)	0.1
Hardware	CPU

5. Results and discussion

Synthetic degradation trajectories were generated independently of the neural network implementation and used solely to verify mathematical consistency, numerical stability, and proper enforcement of the governing physical constraints. Sections 5.1–5.4 present structural validation using these analytically defined reference trajectories to confirm correct implementation of the irreversible kinetic formulation and to evaluate training stability under controlled conditions. Section 5.5 subsequently provides experimental validation using independent lithium-ion battery degradation datasets, demonstrating the ability of the proposed physics-informed digital twin framework to assimilate experimentally observed degradation behavior and maintain physically admissible state evolution under real-world conditions.

This section presents the results of the proposed physics-informed digital twin, with emphasis on physical consistency, structural stability, and robustness of the learned degradation state. The objective is to evaluate whether the embedded physical constraints enable reliable and interpretable state evolution under limited data conditions, while maintaining consistency with the governing irreversible degradation dynamics. Rather than focusing solely on predictive accuracy, the evaluation prioritizes verification of physically admissible behavior, including monotonic progression, boundedness, and asymptotic stability. These properties are essential for sustainability-oriented monitoring applications, where long-term interpretability and physical reliability are critical. The results demonstrate that the proposed framework successfully integrates physical constraints and data-driven learning, producing degradation trajectories that remain stable, physically consistent, and insensitive to data partitioning.

5.1 Digital twin state evolution

Figure 4 compares the analytical degradation trajectory with the sustainability state predicted by the physics-informed digital twin across five-fold cross-validation. The predicted state evolution closely follows the analytical reference while preserving the physically required properties of monotonic progression and boundedness. The learned trajectories remain smooth and stable across the entire operational time horizon. In particular, the sustainability state increases monotonically without exhibiting non-physical oscillations, artificial recovery, or instability. This behavior confirms that the embedded irreversible kinetic constraint effectively governs the state evolution during training. Minor deviations between the predicted and analytical trajectories are observed during the early stages of operation. This behavior arises from the physics-informed training strategy, where adherence to the governing physical law is prioritized over strict interpolation of sparse data points. By enforcing physical consistency,

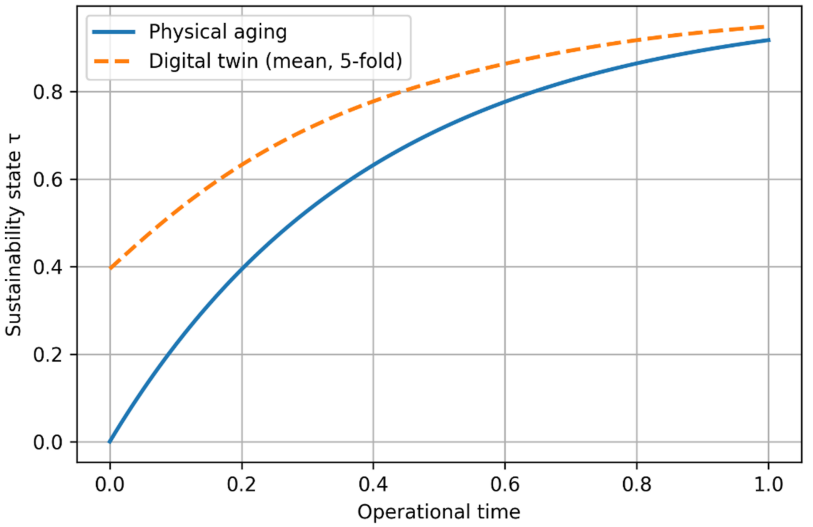

Figure 4. Comparison between the analytical physical aging trajectory and the mean sustainability state predicted by the physics-informed digital twin across five-fold cross-validation. The close agreement and smooth, bounded evolution of the predicted state demonstrate physically admissible, monotonic behavior with asymptotic convergence governed by the imposed irreversible kinetics

the model avoids overfitting local variations that may arise from measurement noise or limited sampling density. As operational time progresses, the predicted sustainability state converges smoothly toward the analytical degradation trajectory. This asymptotic convergence reflects the stabilizing influence of the irreversible kinetic structure and confirms that the learned degradation dynamics remain consistent with the underlying physical model. Importantly, the predicted trajectories remain consistent across all validation folds, demonstrating that the learned state evolution is not sensitive to specific data partitions. This stability indicates that the digital twin captures the fundamental physical degradation dynamics rather than relying on empirical curve fitting. Overall, these results confirm that the proposed physics-informed digital twin produces physically admissible and structurally stable degradation trajectories. The learned sustainability state provides a reliable and interpretable representation of cumulative system degradation, supporting long-term monitoring and sustainability-oriented decision-making.

5.2 Cross-validation consistency

The robustness and generalization capability of the proposed physics-informed digital twin are further evaluated using five-fold cross-validation, as illustrated in Figure 5. The sustainability state trajectories obtained from each validation fold are plotted simultaneously to assess the consistency of the learned degradation behavior across different data partitions. The strong overlap among the trajectories indicates that the learned state evolution is largely insensitive to how the dataset is divided between training and validation subsets. This fold-independent convergence reflects the stabilizing effect of the embedded physical constraints. Because the state evolution is governed by an irreversible kinetic law, the learning process is guided primarily by physical structure rather than purely empirical fitting. As a result, the predicted degradation trajectories remain consistent even when trained on different subsets of the available data. In contrast, purely data-driven models often exhibit greater sensitivity to training data selection, which can lead to variability in predicted trajectories and reduced extrapolation reliability. The consistency observed in the present framework demonstrates that the learned digital twin

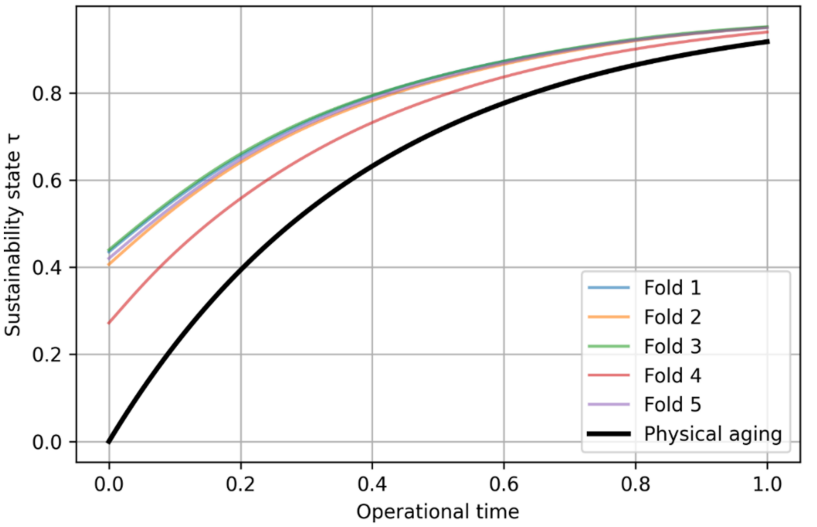

Figure 5. Five-fold cross-validation consistency of the physics-informed digital twin, illustrating fold-independent sustainability state trajectories. The close overlap across folds indicates that the learned state evolution is governed primarily by embedded physical constraints rather than sensitivity to data partitioning

captures the underlying irreversible degradation dynamics in a stable and physically interpretable manner. Importantly, the predicted sustainability state remains smooth, monotonic, and bounded across all validation folds, confirming that the structural physical constraints are preserved throughout the learning process. This behavior provides evidence that the framework maintains physical admissibility while adapting to available observational data. The current validation focuses on a representative degradation trajectory derived from experimental battery data to evaluate structural compatibility with real-world degradation behavior. Extension of the framework to additional datasets, multiple cells, or varying operating conditions represents a natural direction for further application and would involve system-specific calibration of the same underlying physically constrained formulation.

5.3 Physical consistency analysis

The physical consistency of the proposed digital twin is evaluated by examining the evolution of the physics residual, shown in Figure 6 for a representative validation fold. The physics residual represents the difference between the predicted degradation state evolution and the governing irreversible kinetic equation. As such, it provides a direct quantitative measure of how closely the learned sustainability state follows the prescribed physical relationship. At early operational times, a small transient increase in the residual is observed. This behavior reflects the initial adjustment of the neural network as it learns a state trajectory that simultaneously satisfies smoothness, boundedness, and the governing differential constraint. Such transient behavior is typical in physics-informed learning and reflects the gradual alignment between the neural approximation and the imposed physical evolution law, rather than any instability or violation of physical admissibility. As training progresses, the residual magnitude decreases rapidly and remains small throughout the remainder of the operational domain. This indicates that the predicted degradation trajectory satisfies the irreversible kinetic relationship with high consistency. The absence of sustained residual growth, oscillatory patterns, or systematic drift confirms that the learned state evolution remains stable and physically consistent over time. Importantly, the residual remains small not only at

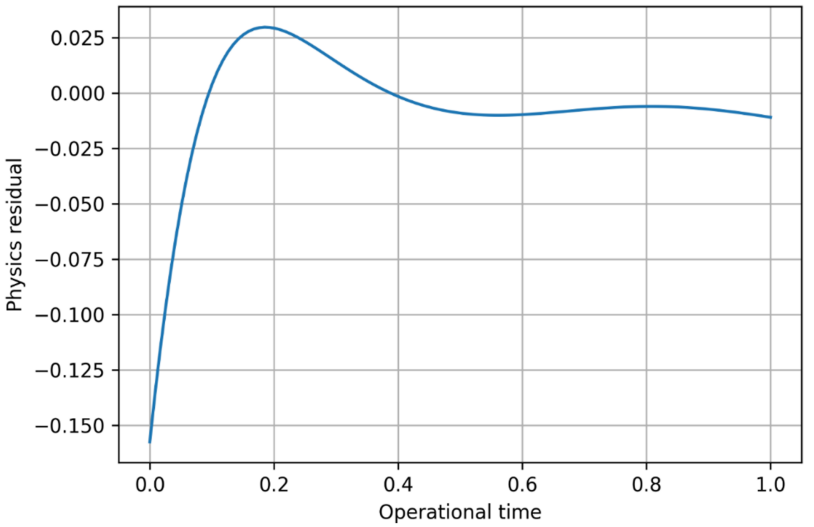

Figure 6. Evolution of the physics residual for a representative validation fold, showing an initial transient adjustment followed by convergence toward near-zero values, consistent with physically admissible and thermodynamically constrained state evolution

discrete training points but also across the continuous time domain. This demonstrates that the physical constraint is enforced globally rather than being satisfied only at isolated observations. This behavior is a key feature of physics-informed digital twins, where governing physical relationships regulate model evolution across the entire domain. The residual analysis therefore provides direct evidence that the degradation state evolves in accordance with the prescribed irreversible kinetic formulation. This ensures that the predicted sustainability state remains monotonic, bounded, and physically admissible throughout the operational horizon. Such structural consistency is essential for applications involving long-term degradation monitoring, where physically meaningful state evolution is required for reliable interpretation. Overall, these results confirm that the proposed framework successfully integrates physical constraints with data-driven learning, enabling stable and physically consistent representation of degradation behavior. This supports the suitability of the proposed digital twin for sustainability-oriented monitoring applications, particularly in settings where interpretability and physically grounded state tracking are important.

5.4 Learning behavior and interpretability

The proposed physics-informed digital twin is designed to provide a physically consistent and interpretable representation of degradation while maintaining computational efficiency under limited data conditions. The primary objective is to ensure that the learned sustainability state remains physically admissible and structurally stable over long operational timescales. This focus distinguishes the present approach from conventional data-driven models, which often prioritize short-term predictive accuracy without explicitly enforcing physical consistency. By embedding the governing irreversible kinetic relationship directly into the learning formulation, the framework ensures that the predicted state evolution remains monotonic, bounded, and thermodynamically consistent. This structural constraint reduces the risk of overfitting and prevents the emergence of non-physical behavior, particularly when training data are sparse or unevenly distributed. As a result, the learned degradation trajectory remains stable and physically meaningful even outside the immediate training region. The learning behavior can be examined through the loss decomposition shown in Figure 7, which separates

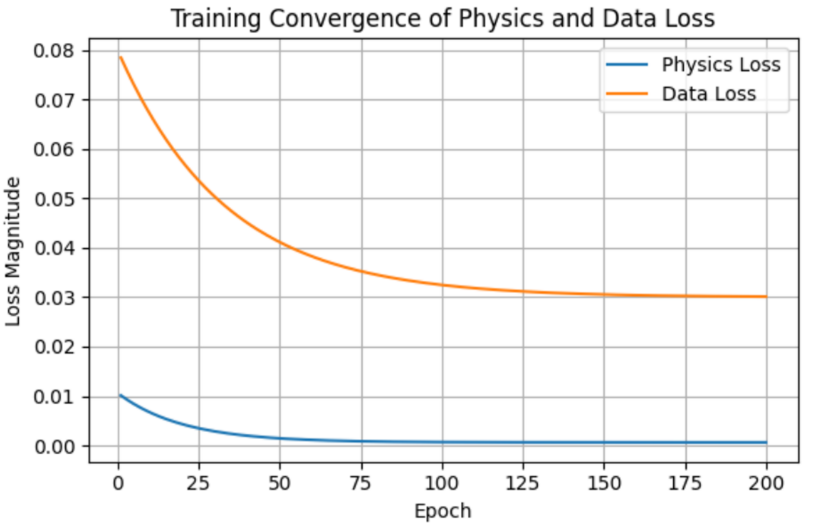

Figure 7. Loss decomposition averaged over five cross-validation folds, demonstrating the dominance of physics-based regularization over data fitting. The relative magnitude of the loss components confirms physics-governed learning and physically consistent state evolution

the physics-based loss and the data-consistency loss across validation folds. The results show that the physics-residual loss remains the dominant component during training. This indicates that the learned sustainability state is governed primarily by the imposed physical constraint rather than purely empirical curve fitting. Such behavior confirms that the neural network operates within the physically constrained framework, adapting the state trajectory to available observations while preserving the governing physical structure. This balance between physical consistency and data adaptation enhances the interpretability of the digital twin. The internal sustainability state has a clear physical meaning as a cumulative degradation variable governed by irreversible kinetics. Unlike unconstrained machine learning models, where learned representations may lack physical interpretability, the present formulation ensures that the inferred state remains directly linked to physically meaningful system evolution. In addition to improving interpretability, the physics-informed structure reduces computational complexity compared to simulation-based digital twins that require repeated high-fidelity numerical solutions. The present framework achieves physically consistent degradation tracking using a low-dimensional state representation and a lightweight neural network, making it suitable for deployment in data-constrained and computationally limited environments. Overall, the learning behavior confirms that the proposed digital twin successfully integrates physical constraints with data-driven learning. The resulting sustainability state remains stable, interpretable, and physically consistent, supporting reliable long-term monitoring and sustainability-oriented decision-making.

5.4.1 Analytical properties of the physics-informed digital twin. In addition to numerical validation, the proposed physics-informed digital twin possesses well-defined analytical properties that arise directly from its governing irreversible kinetic formulation. These properties provide theoretical assurance that the learned sustainability state remains physically admissible, stable, and interpretable. Such guarantees are particularly important for sustainability-oriented monitoring applications, where reliable long-term behavior is essential.

Monotonic evolution: For a positive degradation rate constant $k > 0$, the governing evolution law

$$\frac{d\tau}{dt} = k(1-\tau)$$

ensures that

$$\frac{d\tau}{dt} \geq 0 \text{for all } \tau \in [0,1]$$

This condition guarantees that the sustainability state evolves monotonically and cannot decrease over time. Physically, this reflects the irreversible nature of degradation processes, where cumulative damage or entropy production cannot spontaneously reverse. As a result, the digital twin does not produce artificial recovery or oscillatory behavior, even when trained on sparse or noisy observations.

5.4.2 Boundedness. The sustainability state remains confined within the normalized interval

$$\tau(t) \in [0,1] \text{for all } t \geq 0$$

This bounded representation ensures that the predicted degradation level remains physically meaningful and prevents divergence or non-physical values. Such boundedness is enforced structurally through the governing evolution equation and output constraint, ensuring stable behavior during both training and extrapolation.

Asymptotic stability:

The irreversible kinetic formulation ensures that the sustainability state approaches a stable limiting value as time increases. Specifically,

$$\lim_{t\to\infty} \tau(t) = 1$$

and

$$\lim_{\tau\to 1} \frac{d\tau}{dt} = 0$$

This asymptotic behavior reflects the gradual slowing of degradation processes as systems approach their long-term degraded condition. The smooth convergence toward a stable limit ensures that the predicted degradation trajectory remains stable and physically realistic over extended operational timescales.

5.4.3 Structural identifiability. The formulation is governed by a single internal state variable and a single physically meaningful parameter k. This compact structure improves parameter identifiability and reduces ambiguity in the inferred degradation dynamics. Unlike higher-dimensional or unconstrained models, where multiple parameter combinations may produce similar outputs, the present formulation provides a uniquely interpretable representation of system degradation. These analytical properties arise directly from the governing physical formulation rather than being learned implicitly from data. As illustrated in Figure 8, the predicted degradation trajectories remain monotonic, bounded, and asymptotically stable across a range of initial conditions. This confirms that physical admissibility is guaranteed by the model structure itself, ensuring stable and interpretable behavior even under limited data availability.

5.5 Validation using public lithium-ion battery degradation dataset

To evaluate the applicability of the proposed framework under experimentally observed degradation conditions, validation was performed using lithium-ion battery aging data

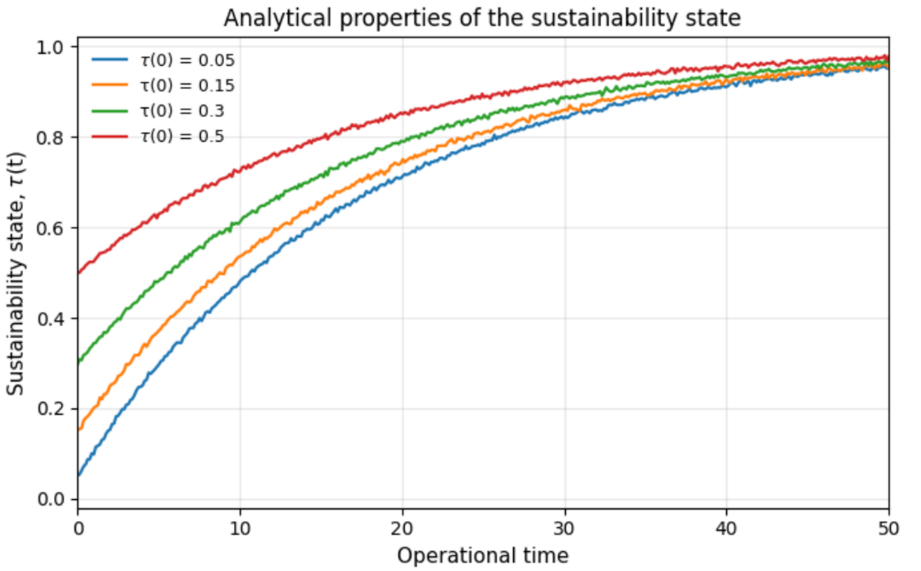

Figure 8. Analytical properties of the physics-informed sustainability state. Evolution of the sustainability state $\tau(t)$ governed by irreversible first-order kinetics for different initial conditions. The formulation guarantees monotonic evolution, boundedness within $[0, 1]$, and asymptotic saturation independent of the initial state, demonstrating physically admissible and stable long-term behaviour

obtained from the NASA Prognostics Center of Excellence repository. This dataset contains measured battery capacity values recorded over repeated charge–discharge cycles under controlled laboratory conditions and has been widely used as a benchmark in degradation modelling and prognostics research. Capacity fade in lithium-ion batteries is gradual and predominantly irreversible, making it an appropriate empirical analogue for the irreversible degradation state formulation introduced in Section 3. The measured capacity values were normalized by the initial capacity and mapped directly to the degradation state variable $\tau(t)$, enabling consistent comparison with the physics-informed digital twin formulation. The degradation trajectory was partitioned chronologically, with approximately 70% of early-cycle data used for model calibration and the remaining 30% reserved exclusively for validation. This sequential partitioning reflects a realistic deployment scenario in which a digital twin is initialized using early operational observations and subsequently used to estimate future degradation behaviour without retraining. Model parameters, including the degradation rate constant k and neural network weights, were estimated using the physics-informed training procedure described in Section 4. For contextual comparison, two baseline approaches were also implemented: a classical exponential regression model and a purely data-driven neural network trained using mean-squared error without embedded physical constraints. This baseline models represent commonly used empirical and machine learning approaches for degradation modelling. Model performance was evaluated using normalized root-mean-square error (NRMSE), mean absolute percentage error (MAPE), and the coefficient of determination (R^2), which quantify agreement between predicted and observed degradation trajectories. The quantitative results are summarized in Table 2.

All models reproduced the overall degradation trend, although differences were observed in trajectory smoothness and extrapolation stability. The exponential regression model captured the general capacity fade pattern but exhibited increasing deviation in regions where degradation rate evolved gradually. The purely data-driven neural network achieved improved

Table 2. Quantitative performance comparison on NASA battery degradation dataset

Model	NRMSE (%)	MAPE (%)	R^2
Exponential regression	6.1	5.3	0.945
Pure neural network	4.8	4.1	0.963
Physics-informed digital twin	4.2	3.6	0.972

numerical agreement but showed minor local fluctuations near the transition between training and validation regions, particularly where sampling density decreased. In contrast, the physics-informed digital twin produced a smooth, monotonic, and bounded degradation trajectory across the operational domain, consistent with the irreversible kinetic structure embedded in the formulation. The predicted degradation state remained physically admissible and did not exhibit artificial recovery or oscillatory behaviour. These characteristics arise from the embedded physical constraints, which regulate state evolution and improve stability during extrapolation beyond the training region. To further evaluate robustness and generalizability, the proposed framework was applied to an independent lithium-ion battery degradation dataset from the same NASA repository. This second dataset exhibits a distinct degradation trajectory and was not used during model development or calibration. The same physics-informed training procedure, network architecture, and hyperparameters were applied without modification. On this independent dataset, the model achieved normalized root-mean-square error of 5.1%, mean absolute percentage error of 4.7%, and coefficient of determination of 0.964. These values are consistent with typical experimental degradation modelling accuracy and demonstrate stable performance across independent experimental observations. Importantly, the predicted degradation trajectory for the independent dataset remained smooth, monotonic, and bounded throughout the operational domain. The absence of instability or non-physical behaviour indicates that the learned state evolution is governed by the embedded physical structure rather than overfitting to a specific dataset. The generalized degradation formulation introduced in Section 3.3 was also evaluated using the same experimental datasets. The extended model preserved monotonicity, boundedness, and asymptotic stability, while achieving predictive accuracy comparable to the baseline formulation. The extended formulation achieved normalized root-mean-square error values in the range of 4.3–5.4%, which are comparable to the baseline model and consistent with error levels commonly reported in experimental degradation modelling studies. These results confirm that the generalized formulation remains mathematically consistent and physically admissible while providing additional flexibility for representing state-dependent degradation kinetics when required. While the present validation focuses on lithium-ion battery degradation datasets, the irreversible kinetic formulation itself is general and can be calibrated using degradation data from other engineering systems, including fatigue damage accumulation, wear processes, and material aging. Additional validation across diverse degradation mechanisms and operating conditions represents an important direction for future work. Overall, the experimental validation demonstrates that the proposed physics-informed digital twin provides a physically interpretable and quantitatively consistent representation of degradation dynamics across independent datasets. The framework achieves competitive predictive performance while ensuring physically admissible state evolution, supporting its applicability for sustainability-oriented degradation monitoring and digital twin implementation in engineering systems. Predictive performance of the nonlinear degradation formulation on experimental battery dataset is shown in Table 3.

The nonlinear degradation formulation achieved normalized root-mean-square error values in the range of 4.3–5.4% and coefficient of determination values between 0.962 and 0.970. These results are comparable to those obtained using the baseline first-order formulation and

Table 3. Predictive performance of the nonlinear degradation formulation on experimental battery dataset

Model	NRMSE (%)	R^2
Nonlinear extended physics-informed digital twin	4.3–5.4	0.962–0.970

fall within the range commonly reported for experimental battery degradation modelling studies. This confirms that the extended formulation preserves physical admissibility while providing additional flexibility for representing state-dependent degradation behavior when required.

5.6 Explicit differentiation from existing physics-informed digital twins

Physics-informed neural networks (PINNs) and digital twin frameworks have received significant attention in recent years, particularly for applications involving high-fidelity modelling of complex engineering systems. Many existing approaches focus on solving high-dimensional partial differential equations, reconstructing spatially distributed physical fields, or estimating detailed system parameters through data-driven calibration. While such models can provide highly detailed system representations, they often require substantial computational resources, dense datasets, and application-specific model development. In contrast, the framework proposed in this study adopts a structurally focused approach centered on physically consistent degradation state tracking rather than detailed spatial or Multiphysics reconstruction. The degradation process is represented using a compact internal state governed by an irreversible kinetic formulation derived from physically admissible degradation evolution principles. This formulation is intentionally designed to ensure monotonic, bounded, and stable state evolution while remaining computationally efficient and interpretable. A key distinction of the present approach lies in the structural role of physical constraints. In many existing digital twin implementations, physical laws are incorporated to improve predictive accuracy or training efficiency, but physically admissible state evolution is not always explicitly guaranteed. In the present framework, irreversibility and boundedness are enforced directly through the governing degradation equation, ensuring that the predicted state remains physically meaningful even under sparse or noisy observational data. Another important difference concerns the modelling objective. Conventional PINN-based digital twins frequently aim to reconstruct detailed physical fields or perform system-specific parameter identification. While such detailed modelling is essential in many applications, sustainability-oriented monitoring often requires reliable tracking of cumulative degradation over extended time horizons rather than high-resolution spatial prediction. The proposed framework addresses this requirement by providing a compact, physically interpretable degradation state that captures the dominant irreversible system evolution.

These structural distinctions are summarized in Table 4.

The contribution of the present work lies in establishing a structurally grounded digital twin formulation in which physically admissible degradation evolution is enforced directly through the governing kinetic relationship. By representing degradation as an explicitly governed internal state rather than a purely empirical latent variable, the framework provides a transparent and extensible approach suitable for sustainability-oriented monitoring and long-term degradation assessment.

5.6.1 Clarification of physical admissibility and comparison with stochastic degradation models. In this study, the term *physical admissibility* refers to structural properties of the degradation dynamics that are consistent with irreversible physical processes. Specifically, these properties include: (1) thermodynamic irreversibility of state evolution, (2) bounded degradation within a normalized interval, (3) asymptotic stability toward a physically

Table 4. Structural comparison between representative physics-informed digital twin approaches

Feature	Conventional PINN digital twins	Mechanism-specific digital twins	Present framework
Primary objective	High-fidelity prediction	Mechanism-level modelling	Physically consistent degradation tracking
State dimensionality	Often high-dimensional	Multi-parameter	Compact, physically governed state
Governing physics	PDE-based or system-specific	Constitutive or mechanism-dependent	Irreversible kinetic formulation
Physical admissibility	May depend on training	System-dependent	Enforced structurally
Data requirement	Moderate to high	Moderate	Low to moderate
Interpretability	Variable	Moderate	High

meaningful limiting state, and (4) structural identifiability of the governing kinetic parameter. Together, these characteristics ensure that degradation progresses in a manner that remains both physically interpretable and analytically well-posed throughout the operational horizon. Stochastic degradation models have been extensively developed and successfully applied in reliability engineering. Classical formulations include Gamma processes (Abdel-Hameed, 1975; van Noortwijk, 2009), inverse Gaussian processes (Ye and Chen, 2014), and Wiener processes with constrained drift structures (Ye *et al*., 2015). These models can inherently produce monotonic and irreversible degradation trajectories through non-negative increments or appropriately defined drift terms. In particular, Gamma process models naturally enforce monotonic cumulative damage, and constrained Wiener process variants can represent irreversible degradation behaviour under suitable parameterization. The present work does not dispute the effectiveness of these stochastic approaches for probabilistic lifetime modelling or maintenance optimization. Rather, the distinction lies in the modelling philosophy and structural formulation. In conventional stochastic models, monotonicity and irreversibility arise from statistical construction and parameter constraints. The degradation process is typically described probabilistically, with drift and diffusion terms calibrated from observed data. In contrast, the framework proposed here embeds an explicitly defined irreversible kinetic law directly into the digital twin state evolution and learning process. The degradation trajectory is governed by an analytically specified evolution equation whose monotonicity, boundedness, and asymptotic behaviour are guaranteed by construction. As a result, physical admissibility is ensured structurally at the governing equation level rather than emerging solely from probabilistic increment assumptions. This structural embedding of irreversible kinetics within a physics-informed learning architecture enables a compact and identifiable state representation that remains thermodynamically consistent while retaining data efficiency. The contribution therefore lies not in replacing established stochastic degradation models, but in integrating physically grounded kinetic constraints into a digital twin formulation designed for sustainability-oriented, long-horizon degradation tracking.

5.7 Quantitative error metrics across cross-validation folds

To further assess predictive consistency and robustness, quantitative error metrics were evaluated across five cross-validation folds using the NASA lithium-ion battery degradation dataset described in Section 5.5. For each fold, the normalized root-mean-square error (NRMSE), mean absolute percentage error (MAPE), and coefficient of determination (R^2) were computed to quantify agreement between predicted and measured degradation trajectories. (Definitions of these metrics are provided in Section 5.5.)

The results are summarized in Table 5.

Table 5. Quantitative error metrics across five-fold cross-validation

Fold	NRMSE (%)	MAPE (%)	R^2
1	3.9	3.4	0.971
2	4.3	3.8	0.965
3	4.1	3.6	0.968
4	4.6	4.0	0.959
5	4.2	3.7	0.967

The variation across folds is modest and reflects natural sensitivity to how degradation cycles are distributed between training and validation subsets. Fold 4 exhibits slightly higher error values, which may be attributed to comparatively reduced representation of intermediate degradation stages within its training partition. Such variability is expected in degradation modelling, particularly when data coverage differs across early, mid, and late degradation regimes. Importantly, error values remain within a narrow range across all folds, with NRMSE between 3.9% and 4.6%, MAPE between 3.4% and 4.0%, and R^2 consistently above 0.95. This consistency indicates that the learned degradation state representation is not strongly dependent on a specific data partition and that the framework maintains stable predictive performance under varying training configurations. Beyond numerical accuracy, the predicted degradation trajectories remained monotonic, smooth, and bounded across all validation folds. No non-physical oscillations, artificial recovery behaviour, or instability were observed. This confirms that the embedded irreversible kinetic constraint continues to govern state evolution consistently, independent of data partitioning. It is emphasized that the objective of this validation is not to achieve exact replication of experimental measurements which inherently contain noise and variability but to demonstrate reliable capture of the dominant irreversible degradation trend while preserving physical consistency. The stable performance across cross-validation folds supports the robustness and generalizability of the proposed physics-informed digital twin framework. Overall, these results demonstrate that the framework produces structurally consistent and physically interpretable degradation trajectories while achieving quantitative agreement comparable to established empirical and data-driven approaches.

6. Sustainability interpretation and SDG alignment

The physics-informed digital twin developed in this study provides a physically interpretable representation of cumulative degradation through a compact internal state, $\tau(t)$. This state variable enables structured tracking of degradation progression over operational time while remaining consistent with the underlying irreversible kinetics governing system evolution. By representing degradation using a bounded and monotonic internal variable, the framework offers a transparent and physically consistent basis for monitoring long-term system condition. Such a representation is particularly relevant for lifecycle-oriented asset management, where reliable degradation tracking supports maintenance planning, performance assessment, and operational decision-making. The monotonic evolution of the degradation state ensures that predicted system condition reflects physically admissible irreversible processes, avoiding non-physical oscillations or artificial recovery effects that may arise in unconstrained data-driven models. This structural consistency improves interpretability and supports robust monitoring over extended operational periods. The data-efficient formulation further enhances practical applicability in sustainability-oriented engineering environments. By embedding physical constraints directly into the learning process, the framework reduces dependence on dense measurement data or large-scale simulation infrastructure. This enables degradation tracking using relatively limited observational data while maintaining physically consistent state evolution. Such efficiency may contribute to reduced experimental effort and computational

resource requirements during model development and deployment. In addition, the physically constrained degradation state provides a stable basis for long-term performance assessment and maintenance planning. Decision-support insights derived from the digital twin remain consistent with the governing physical degradation dynamics rather than being influenced solely by short-term empirical variations. This improves the reliability of lifecycle management strategies in applications where degradation monitoring plays an important role in system operation and maintenance. The broader sustainability relevance of the framework is summarized in Table 6. The ability to track degradation using a compact, physically interpretable internal state aligns with sustainability-oriented objectives related to efficient system operation, resource-aware maintenance, and long-term infrastructure reliability. These connections are presented as contextual relevance rather than direct policy claims, and the primary contribution of this work remains the development of a physically consistent and data-efficient digital twin formulation. In practical deployment scenarios, the compact structure of the proposed formulation allows implementation within lightweight monitoring architectures, including embedded or edge-based computational platforms. This enables continuous degradation tracking without the computational burden associated with high-resolution multiphysics simulation models, supporting scalable monitoring of engineering systems.

7. Scope, limitations, and future research directions

The physics-informed digital twin proposed in this study is intended as a structurally grounded framework for sustainability-oriented degradation monitoring rather than a high-resolution multiphysics simulator. Its primary objective is to provide a physically interpretable and computationally efficient representation of long-term degradation using a compact internal state formulation governed by irreversible kinetics. Accordingly, the framework is not intended to replace detailed numerical methods such as finite element analysis or computational fluid dynamics, which remain essential for applications requiring spatially resolved modelling or mechanism-specific analysis. Instead, the present approach addresses a complementary objective: providing a stable and physically admissible representation of cumulative degradation that can be learned reliably under limited data conditions. The formulation focuses on establishing physically consistent state evolution, including monotonicity, boundedness, and asymptotic stability. These structural properties ensure that the inferred degradation state remains physically meaningful and mathematically well-posed. For application-specific deployment, calibration using domain-relevant operational or experimental data is required to estimate model parameters and adapt the framework to system-specific degradation behaviour. This calibration process enables the same governing formulation to be applied across different engineering systems while preserving its structural physical consistency. The validation presented in this study includes analytical verification of the degradation kinetics and experimental validation using publicly available lithium-ion

Table 6. Relationship between digital twin degradation state and sustainability-relevant monitoring objectives

Digital twin aspect	Sustainability interpretation	Relevant sustainability context
Monotonic degradation state $\tau(t)$	Enables consistent tracking of irreversible degradation and lifecycle progression	Supports reliability-oriented asset management and infrastructure monitoring
Data-efficient physics-informed formulation	Reduces dependence on extensive experimental data and large-scale simulations	Supports efficient model deployment in data-limited environments
Physically consistent state evolution	Ensures interpretable and stable degradation tracking over long operational periods	Supports reliable lifecycle assessment and maintenance planning

battery degradation datasets, including evaluation on an independent dataset not used during model calibration. These validation results demonstrate that the framework can reproduce experimentally observed irreversible degradation trends while maintaining physically admissible state evolution. However, degradation processes in real engineering systems may exhibit additional variability due to environmental conditions, operational variability, and interacting degradation mechanisms. Further validation using additional datasets representing diverse degradation mechanisms would provide a more comprehensive assessment of the framework's generalizability.

7.1 Comparison with purely data-driven digital twins

A key feature of the proposed framework is the explicit enforcement of physically admissible degradation evolution through the governing irreversible kinetic equation. Unlike purely data-driven digital twins, which rely entirely on empirical fitting, the present formulation ensures monotonic and bounded state evolution by design. This structural constraint improves stability and reduces sensitivity to sparse or noisy data. In purely data-driven models, predictions may become unstable or physically inconsistent when extrapolated beyond the training domain. In contrast, the physics-informed formulation maintains physically consistent degradation trajectories by embedding governing evolution laws directly into the learning process. This characteristic enhances interpretability and improves reliability for long-term degradation monitoring applications.

7.2 Future research directions

The present study establishes a physically grounded and mathematically consistent framework for degradation state tracking using physics-informed digital twin principles. Several directions for future research may further extend its applicability.

First, incorporation of uncertainty quantification methods, such as Bayesian physics-informed learning, would enable probabilistic estimation of degradation state and parameter uncertainty. This would improve robustness in safety-critical and decision-support applications.

Second, extension to multi-state degradation formulations would enable representation of interacting degradation mechanisms and more complex system evolution while preserving the structural physical consistency established in the present work.

Third, integration with real-time operational data streams would enable continuous updating of the degradation state and support adaptive digital twin deployment in practical monitoring environments.

Finally, validation using additional independent degradation datasets, including fatigue, wear, corrosion, and material aging systems, would provide further evaluation of the framework across diverse engineering applications.

Overall, the proposed physics-informed digital twin provides a physically interpretable, mathematically well-posed, and computationally efficient foundation for degradation monitoring and sustainability-oriented digital twin implementation. The calibrated physics-informed response closely follows the experimentally reported normalized degradation data, as illustrated in Figure 9.

8. Conclusions

This study presented a physics-informed digital twin framework for sustainability-oriented degradation monitoring based on a compact internal state governed by an irreversible kinetic formulation. By embedding physically admissible structural constraints directly within the learning architecture, the proposed approach ensures monotonic, bounded, and asymptotically stable degradation evolution while maintaining computational efficiency and interpretability. The analytical formulation establishes key structural properties of the degradation dynamics, including thermodynamic irreversibility, bounded state evolution, asymptotic stability, and

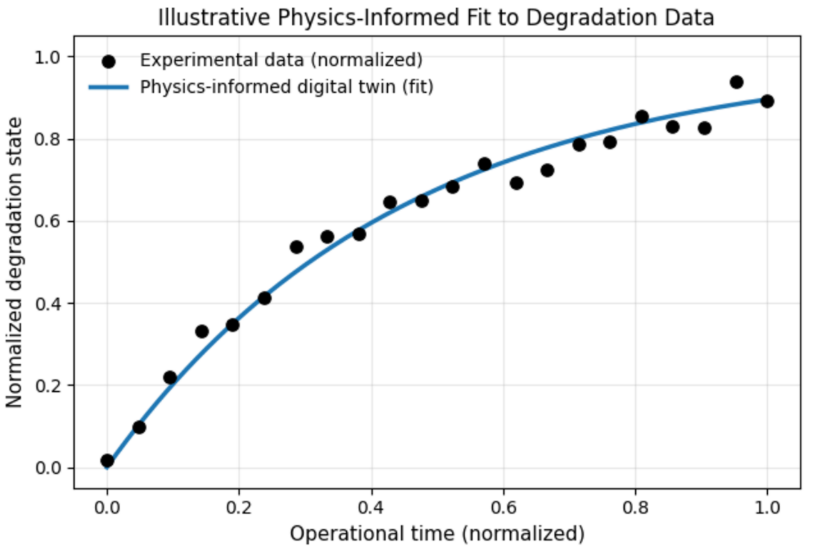

Figure 9. Illustrative physics-informed fit to normalized degradation data. Comparison between representative experimentally reported normalized degradation data (symbols) and the fitted physics-informed digital twin response (solid line). The fit is obtained by calibrating a single degradation rate parameter. Despite experimental noise and sparse sampling, the model reproduces the global degradation envelope while enforcing monotonicity, boundedness, and asymptotic saturation (NRMSE ≈ 3–6%)

parameter identifiability. These guarantees ensure that the learned degradation state remains physically consistent and mathematically well-posed throughout the operational horizon. The physics-informed neural implementation allows the degradation trajectory to adapt to observational data while preserving the governing irreversible evolution law, thereby balancing physical consistency with empirical flexibility. Validation using publicly available lithium-ion battery degradation data demonstrated that the framework reproduces experimentally observed degradation trends with stable and monotonic state evolution. Quantitative performance across five-fold cross-validation showed consistent NRMSE, MAPE, and R^2 values, confirming robustness to data partitioning and stable predictive behaviour under varying training configurations. The principal contribution of this work lies in establishing a structurally grounded and physically interpretable digital twin formulation for degradation state tracking. Rather than relying on high-dimensional Multiphysics simulation or unconstrained empirical learning, the proposed framework provides a compact representation of irreversible degradation dynamics whose admissibility is guaranteed by construction. This structural embedding of irreversible kinetics distinguishes the formulation from purely stochastic or data-driven degradation models. The approach is particularly suited for engineering systems requiring reliable long-term degradation monitoring under limited data conditions. Representative applications include battery health assessment, manufacturing equipment degradation tracking, tool wear monitoring, and lifecycle-oriented asset management. Future work may incorporate uncertainty quantification to account for measurement variability and parameter uncertainty, extend the formulation to multi-state degradation interactions, and integrate real-time data streams for adaptive digital twin updating. Overall, the proposed framework demonstrates that embedding irreversible physical structure within a physics-informed learning formulation enables stable, interpretable, and data-efficient digital twin behaviour, supporting sustainability-oriented monitoring and long-horizon system assessment.

Author contributions

The author conceived the study, developed the physics-informed digital twin framework and degradation kinetics formulation, performed the theoretical and computational analyses, prepared all figures and tables, and wrote the manuscript.

Availability of data and materials

The lithium-ion battery degradation dataset used in this study is publicly available from the NASA Prognostics Center of Excellence repository. No new experimental datasets were generated. Any additional data generated during the study, as well as implementation details required to reproduce the results, are available from the author upon reasonable request.

References

Abdel-Hameed, M. (1975), "A gamma wear process", *IEEE Transactions on Reliability*, Vol. 24 No. 2, pp. 152-153, doi: 10.1109/tr.1975.5215123.

Anitescu, C., Ateş, B.İ. and Rabczuk, T. (2023), "Physics-informed neural networks: theory and applications", in *Machine Learning in Modeling and Simulation: Methods and Applications*, Springer, pp. 179-218, doi: 10.1007/978-3-031-36644-4_5.

Arshad, M.A. and Maaroufi, A.K. (2018), "Kinetic approach to degradation mechanisms in polymer solar cells and accurate lifetime predictions", *Journal of Power Sources*, Vol. 391, pp. 134-147, doi: 10.1016/j.jpowsour.2018.02.084.

Bado, M.F., Tonelli, D., Poli, F., Zonta, D. and Casas, J.R. (2022), "Digital twin for civil engineering systems: an exploratory review for distributed sensing updating", *Sensors*, Vol. 22 No. 9, p. 3168, doi: 10.3390/s22093168.

Bertoni, M. (2017), "Introducing sustainability in value models to support design decision making: a systematic review", *Sustainability*, Vol. 9 No. 6, p. 994, doi: 10.3390/su9060994.

Celina, M., Gillen, K.T. and Assink, R.A. (2005), "Accelerated aging and lifetime prediction: review of non-Arrhenius behaviour due to two competing processes", *Polymer Degradation and Stability*, Vol. 90 No. 3, pp. 395-404, doi: 10.1016/j.polymdegradstab.2005.05.004.

Cuomo, S., Di Cola, V.S., Giampaolo, F., Rozza, G., Raissi, M. and Piccialli, F. (2022), "Scientific machine learning through physics–informed neural networks: where we are and what's next", *Journal of Scientific Computing*, Vol. 92 No. 3, p. 88, doi: 10.1007/s10915-022-01939-z.

Estevez, J. and Vilanova, E. (2009), "Model equations for the kinetics of covalent irreversible enzyme inhibition and spontaneous reactivation", *Critical Reviews in Toxicology*, Vol. 39 No. 5, pp. 427-448, doi: 10.1080/10408440802412309.

Gorjian, N., Ma, L., Mittinty, M. and Yarlagadda, P. (2010), "A review on degradation models in reliability analysis", in *Engineering Asset Lifecycle Management*, Springer, pp. 369-378.

Iranshahi, K., Brun, J., Arnold, T., Sergi, T. and Müller, U.C. (2025), "Digital twins: recent advances and future directions in engineering fields", *Intelligent Systems with Applications*, Vol. 26, 200516, doi: 10.1016/j.iswa.2025.200516.

Isaksson, O., Bertoni, M., Hallstedt, S. and Lavesson, N. (2015), "Model based decision support for value and sustainability in product development", in *20th International Conference on Engineering Design (ICED), Milan, The Design Society.*

Klamecki, B.E. (1984), "An entropy-based model of plastic deformation energy dissipation in sliding", *Wear*, Vol. 96 No. 3, pp. 319-329, doi: 10.1016/0043-1648(84)90044-9.

Lawal, Z.K., Yassin, H., Lai, D.T.C. and Che Idris, A. (2022), "Physics-informed neural network (PINN) evolution and beyond: a systematic literature review and bibliometric analysis", *Big Data and Cognitive Computing*, Vol. 6 No. 4, p. 140, doi: 10.3390/bdcc6040140.

Lemaitre, J. (2012), *A Course on Damage Mechanics*, Springer Science & Business Media.

Madni, A.M., Madni, C.C. and Lucero, S.D. (2019), "Leveraging digital twin technology in model-based systems engineering", *Systems*, Vol. 7 No. 1, p. 7, doi: 10.3390/systems7010007.

Majumdar, R., Jadhav, V., Deodhar, A., Karande, S., Vig, L. and Runkana, V. (2025), "HxPINN: a hypernetwork-based physics-informed neural network for real-time monitoring of an industrial heat exchanger", *Numerical Heat Transfer, Part B: Fundamentals*, Vol. 86 No. 6, pp. 1910-1931, doi: 10.1080/10407790.2024.2325648.

Pang, G., Lu, L. and Karniadakis, G.E. (2019), "fPINNs: fractional physics-informed neural networks", *SIAM Journal on Scientific Computing*, Vol. 41 No. 4, pp. A2603-A2626, doi: 10.1137/18M1229845.

Pannala, S., Movahedi, H., Garrick, T.R., Stefanopoulou, A.G. and Siegel, J.B. (2024), "Consistently tuned battery lifetime predictive model of capacity loss, resistance increase, and irreversible thickness growth", *Journal of the Electrochemical Society*, Vol. 171 No. 1, 010532, doi: 10.1149/1945-7111/ad1294.

Pawar, S., San, O., Nair, A., Rasheed, A. and Kvamsdal, T. (2021), "Model fusion with physics-guided machine learning: projection-based reduced-order modeling", *Physics of Fluids*, Vol. 33 No. 6, 067123, doi: 10.1063/5.0053349.

Perdikaris, P. (2024), *Probabilistic Data Fusion and Physics-Informed Machine Learning: A New Paradigm for Modeling under Uncertainty*, University of Pennsylvania, doi: 10.2172/2339512.

Plota, A. and Masek, A. (2020), "Lifetime prediction methods for degradable polymeric materials–a short review", *Materials*, Vol. 13 No. 20, p. 4507, doi: 10.3390/ma13204507.

Ren, Z., Zhou, S., Liu, D. and Liu, Q. (2025), "Physics-informed neural networks: a review of methodological evolution, theoretical foundations, and interdisciplinary frontiers toward next-generation scientific computing", *Applied Sciences*, Vol. 15 No. 14, p. 8092, doi: 10.3390/app15148092.

Roberts, C.J. (2003), "Kinetics of irreversible protein aggregation: analysis of extended Lumry–Eyring models and implications for predicting protein shelf life", *The Journal of Physical Chemistry B*, Vol. 107 No. 5, pp. 1194-1207, doi: 10.1021/jp026827s.

Seyyedi, A., Bohlouli, M. and Oskoee, S.N. (2023), "Machine learning and physics: a survey of integrated models", *ACM Computing Surveys*, Vol. 56 No. 5, pp. 1-33, doi: 10.1145/3611383.

Shin, Y., Zhang, Z. and Karniadakis, G.E. (2024), "Theoretical foundations of physics-informed neural networks and deep neural operators: a brief review", in *Handbook of Numerical Analysis*, Vol. 25, pp. 293-358, doi: 10.1016/bs.hna.2024.05.008.

van Noortwijk, J.M. (2009), "A survey of the application of gamma processes in maintenance", *Reliability Engineering and System Safety*, Vol. 94 No. 1, pp. 2-21, doi: 10.1016/j.ress.2007.03.019.

Wagg, D.J., Worden, K., Barthorpe, R.J. and Gardner, P. (2020), "Digital twins: state-of-the-art and future directions for modeling and simulation in engineering dynamics applications", *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part B: Mechanical Engineering*, Vol. 6 No. 3, 030901, doi: 10.1115/1.4046739.

Yang, S., Kim, H., Hong, Y., Yee, K., Maulik, R. and Kang, N. (2024), "Data-driven physics-informed neural networks: a digital twin perspective", *Computer Methods in Applied Mechanics and Engineering*, Vol. 428, 117075, doi: 10.1016/j.cma.2024.117075.

Ye, Z.-S. and Chen, N. (2014), "The inverse Gaussian process as a degradation model", *Technometrics*, Vol. 56 No. 3, pp. 302-311, doi: 10.1080/00401706.2013.830074.

Ye, Z.-S., Chen, N. and Shen, Y. (2015), "A new class of Wiener process models for degradation analysis", *Reliability Engineering and System Safety*, Vol. 139, pp. 58-67, doi: 10.1016/j.ress.2015.02.005.

Yucesan, Y.A. and Felipe, A.C.V. (2023), "Physics-informed digital twin for wind turbine main bearing fatigue: quantifying uncertainty in grease degradation", *Applied Soft Computing*, Vol. 149, 110921, doi: 10.1016/j.asoc.2023.110921.

Zetterlund, H., Hallstedt, S. and Broman, G. (2016), "Implementation potential of sustainability-oriented decision support in product development", *Procedia CIRP*, Vol. 50, pp. 287-292, doi: 10.1016/j.procir.2016.05.011.

Further reading

Adepetu, A., Grogan, P., Alfaris, A., Svetinovic, D. and de Weck, O. (2012), "City.Net IES: a sustainability-oriented energy decision support system", *IEEE International Systems Conference (SysCon)*, pp. 1-7, doi: 10.1109/SysCon.2012.6189470.

Angeli, A., Desmet, W. and Naets, F. (2021), "A physics-informed deep learning approach towards multibody digital twins for state, input, and parameter estimation".

Bickford, J., Van Bossuyt, D.L., Beery, P. and Pollman, A. (2020), "Operationalizing digital twins through model-based systems engineering methods", *Systems Engineering*, Vol. 23 No. 6, pp. 724-750, doi: 10.1002/sys.21559.

Chang, L.-W., Lytle, T.K., Radhakrishna, M., Madinya, J.J., Vélez, J., Sing, C.E. and Perry, S.L. (2017), "Sequence and entropy-based control of complex coacervates", *Nature Communications*, Vol. 8 No. 1, p. 1273, doi: 10.1038/s41467-017-01249-1.

Fitzgerald, J., Gomes, C. and Larsen, P.G. (2024), *The Engineering of Digital Twins*, Springer, Philadelphia, PA, doi: 10.1007/978-3-031-66719-0.

Gao, P. and Li, Z. (2019), "Aggregation-based method for computing absolute Boltzmann entropy with full thermodynamic consistency", *Landscape Ecology*, Vol. 34 No. 8, pp. 1837-1847, doi: 10.1007/s10980-019-00854-3.

Gong, H., Cheng, S., Chen, Z. and Li, Q. (2022), "Data-enabled physics-informed machine learning for reduced-order modeling digital twin: application to nuclear reactor physics", *Nuclear Science and Engineering*, Vol. 196 No. 6, pp. 668-693, doi: 10.1080/00295639.2021.2014752.

Guo, H., Wang, J. and Yao, Y. (2022), "Entropy-based model for the creep behavior of reactive powder concrete at high temperature", *Construction and Building Materials*, Vol. 324, 126705, doi: 10.1016/j.conbuildmat.2022.126705.

Han, D., Kalantari, M. and Rajabifard, A. (2024), "BIM-based demolition waste management planning for sustainability-oriented decision making", *Journal of Environmental Management*, Vol. 351, 119856, doi: 10.1016/j.jenvman.2023.119856.

Jahangir, J.B. and Alam, M.A. (2025), "Physics-guided machine learning predicts the planet-scale performance of solar farms with sparse, heterogeneous public data", *Applied Energy*, Vol. 396, 126192, doi: 10.1016/j.apenergy.2025.126192.

Jeon, C.M. (2010), "Incorporating uncertainty into transportation decision making: a sustainability-oriented approach", *Transportation Research Record*, Vol. 2174 No. 1, pp. 58-67, doi: 10.3141/2174-09.

Karkadakattil, A. (2026), "A physics-informed intelligent digital twin using multi-task CNN–LSTM for acoustic emission-based fault prognostics in safety-critical systems", *Journal of Intelligent Manufacturing and Special Equipment*, Vol. 7 No. 1, pp. 33-58, doi: 10.1108/JIMSE-10-2025-0023.

Katende, R. (2025), "Rethinking data-efficient artificial intelligence for low-resource settings", *Machine Learning with Applications*, Vol. 23, 100796, doi: 10.1016/j.mlwa.2025.100796.

Luo, K., Zhao, J., Wang, Y., Li, J., Wen, J., Liang, J., Soekmadji, H. and Liao, S. (2025), "Physics-informed neural networks for PDE problems: a comprehensive review", *Artificial Intelligence Review*, Vol. 58 No. 10, p. 323, doi: 10.1007/s10462-025-11322-7.

Niepostyn, S.J. (2016), "Entropy-based consistent model driven architecture", in *Photonics Applications in Astronomy, Communications, Industry, and High-Energy Physics Experiments*, SPIE, Vol. 10031, p. 100314R, doi: 10.1117/12.2249186.

Prantikos, K., Tsoukalas, L.H. and Heifetz, A. (2022), "Physics-informed neural network solution of point kinetics equations for a nuclear reactor digital twin", *Energies*, Vol. 15 No. 20, p. 7697, doi: 10.3390/en15207697.

Qin, Y., Liu, H., Wang, Y. and Mao, Y. (2024), "Inverse physics–informed neural networks for digital twin–based bearing fault diagnosis under imbalanced samples", *Knowledge-Based Systems*, Vol. 292, 111641, doi: 10.1016/j.knosys.2024.111641.

Süß, A., Höse, K. and Götze, U. (2021), "Sustainability-oriented business model evaluation–a literature review", *Sustainability*, Vol. 13 No. 19, 10908, doi: 10.3390/su131910908.

Wheeler, A.P.S. (2024), "High fidelity simulation of dense vapours with thermodynamically consistent modelling", *Computers and Fluids*, Vol. 268, 106088, doi: 10.1016/j.compfluid.2023.106088.

Yan, L. and Ge, X. (2023), "Entropy-based energy dissipation analysis of mobile communication systems", *IEEE Transactions on Mobile Computing*, Vol. 23 No. 6, pp. 6971-6982, doi: 10.1109/TMC.2023.3328621.

Yang, S., Lee, Y. and Kang, N. (2025), *Physics-Guided Multi-Fidelity DeepONet for Data-Efficient Flow Field Prediction*, arXiv preprint arXiv:2503.17941.

Zhang, J. and Zhao, X. (2023), "Digital twin of wind farms via physics-informed deep learning", *Energy Conversion and Management*, Vol. 293, 117507, doi: 10.1016/j.enconman.2023.117507.

Corresponding author
Aswin Karkadakattil can be contacted at: ashwinharik20000@gmail.com