Gearbox Bearing Crack Growth Prognostics with Physics-Informed Machine Learning and Uncertainty Quantification

Mario De Florio[1], Gabriel Appleby[1], Jonathan Keller[1], Ali Eftekhari Milani[2], Donatella Zappalá[2], and Shawn Sheng[1]

[1]National Renewable Energy Laboratory, Golden, CO, 80401, USA
[2]TU Delft, Kluyverweg 1, Delft, 2629 HS, Netherlands

Corresponding authors: Mario De Florio (mario.deflorio@nrel.gov) and Shawn Sheng (shwan.sheng@nrel.gov)

Abstract. This paper introduces the eXtreme Theory of Functional Connections (X-TFC), a physics-informed machine learning algorithm, and tailors it to estimate the remaining useful life (RUL) of wind turbine gearbox bearings experiencing fatigue

5   crack growth. Unlike purely data-driven methods, X-TFC embeds a physics model, based on the Head's theory in this work, into its training objective. The core of X-TFC is a random-projection single-layer neural network trained via Extreme Learning Machine, which requires only limited damage progression data and solves for output weights with a least-squares optimization algorithm. A composite loss function balances the network's fit to observed degradation data against the residuals of the governing crack-growth differential equation, ensuring the learned damage trajectory remains physically plausible. When applied

10   to a vibration-based health-index (HI) dataset measured during the growth of a crack on the inner ring of a high-speed bearing in a wind turbine gearbox (Bechhoefer and Dubé, 2020), X-TFC achieves near-zero prediction bias. Even when trained on only the first 10–20% of the damage progression data, its predictions remain monotonic and smooth, delivering high prognosability and trendability. To quantify the epistemic uncertainty, we employ a Monte Carlo ensemble of independently initialized X-TFC models trained on noise-perturbed data, which yields confidence intervals around each RUL estimate. This approach

15   provides confidence intervals around each RUL estimate, capturing both model-parameter and epistemic uncertainty. In addition to a vibration-based HI, we demonstrate that the proposed framework can be directly applied to a SCADA data-based HI (Eftekhari Milani et al., 2025) measured during similar wind turbine gearbox bearing crack faults, preserving its accuracy and interpretability. This extension shows the versatility of our approach, which is applicable to bearings of multiple gearbox manufacturers, models and ratings using only SCADA data. By integrating domain knowledge with machine learning, X-TFC

20   offers a rapid, reliable tool for crack prognostics. Its adaptability to other bearing failure modes, such as pitch-bearing ring cracks, positions X-TFC as a powerful enabler of data-driven, physics-informed asset management in the wind energy sector and beyond.

1 Introduction

The wind industry has witnessed tremendous growth over the past two decades, as reported by Global Wind Energy Council,
with the global cumulative installed capacity recently crossing the 1,000 gigawatt mark. Although the levelized cost of energy
from wind and wind plant operations and maintenance (O&M) costs have fallen (Wiser et al., 2024); O&M still accounts for
up to 35% of the levelized cost of energy. Approximately half of the wind plant O&M costs are from wind turbine O&M, a
cost of $22/kW-year for land-based wind plants (Wiser et al., 2019). Turbine O&M costs, therefore, represent the single largest
component of wind plant O&M costs and the primary source of potential O&M cost reductions. More reliable components and
better O&M strategies, including for wind turbine blades and drivetrain components, have the potential to reduce premature
failures, O&M costs, downtime, and thus the levelized cost of energy by 10% to 15% (Stehly and Duffy, 2021). Additionally,
an accurate reliability forecast provides crucial information for reducing O&M costs through design improvements, optimized
operation strategies, and enhanced budgeting. Optimized operation strategies often include condition-based maintenance and
prognostics, which rely on the prediction of the remaining useful life (RUL) of components to provide an early warning of
failure and thereby allow maintenance to be scheduled proactively (Bechhoefer et al., 2021).

One challenging reliability problem that remains for wind turbine main bearings and bearings in gearboxes is axial cracking
of the inner rings, examples of which are shown in Figures 1(a) and (b) (Greco et al., 2022). In gearboxes, this failure mode
in intermediate- and high-speed stage bearings accounts for a majority of uptower bearing replacements (Haus et al., 2024).
"Axial" describes the orientation of these cracks as they align with the axis of rotation of the shaft on which the bearing
is mounted. This type of crack is also commonly called a white-etching crack (WEC), due to the appearance of the steel
microstructure when the cracked bearing cross sections are polished, etched with chemicals, and examined under reflected
light. These cracks tend to propagate to spalls or lead to a complete splitting of the bearing inner ring in gearboxes in as little
as 5% to 20% of the bearing rating life (Greco et al., 2013). The conditions leading to WECs, the reasons for their apparent
prevalence in wind turbine main bearings and gearboxes (Demas et al., 2023), and the effectiveness of countermeasures (Jensen
et al., 2021) are now better understood through over a decade of research; however, accurate methods to predict the RUL of
bearings with WECs still need to be developed. Early efforts focused on development of a physics-based probability of failure
model based on frictional energy accumulation, which has been partially validated on a limited set of damaged bearings from
a wind plant; however, the probability of failure and RUL was only predicted only for the population of gearboxes rather than
individual ones (Guo et al., 2020). It was also used to explore the influence of bearing clearance and lubricant temperature
(Clark et al., 2023). Another challenging reliability problem that has recently become more prevalent is fatigue cracking of
pitch bearing outer rings in larger wind turbines as shown in Figure 1(c), resulting in a replacement rate of several percent per
year and a 10% replacement rate for the population in just 5 years (Haus et al., 2025). These cracks can occur at at bolt holes, fill
holes, or locations of significant stiffness changes (Shapiro, 2017) and risk the blade detaching from the turbine. A typical pitch
bearing replacement requires the use of a large, expensive crane and removal of the blade(s) or rotor, with a significant amount
of downtime. However, research into the root causes of pitch bearing ring cracking has only begun (Schleich, 2025; Keller
et al., 2025). In rotating machinery, fatigue cracks in bearings, gears, or shafts are a common failure mode. Classical fracture

Figure 1. Examples of bearing cracks and damage in wind turbines. (a) Gearbox bearing WEC and spalling (Greco et al., 2022). (b) Main bearing WEC and spalling (Greco et al., 2022). (c) Pitch bearing ring crack (Shapiro, 2017).

mechanics models, such as linear elastic fracture mechanics, Head's theory, and dislocation theory, describe such crack growth under cyclic loading. However, no single crack-growth law fits all phases of damage: Bechhoefer and Dubé (2020) showed that three high-cycle fatigue modes were needed to cover the primary crack propagation period. They found that a hybrid

60   combination of modes was more accurate than any individual model.

In recent years, new computational techniques have been developed to improve evaluation and prediction tasks in various fields of science. Among these, Physics-Informed Machine Learning (PIML) (Karniadakis et al., 2021) has emerged as a powerful paradigm for integrating data-driven modeling with known physical laws, often expressed using differential equations. By embedding governing equations and domain constraints directly into learning objectives, PIML methods improve robustness,

65   data efficiency, and interpretability compared to purely data-driven approaches. One of the most widely used PIML frameworks is the Physics-Informed Neural Network (PINN) (Raissi et al., 2019), which augments standard neural-network training with residual losses for differential equations, initial/boundary conditions, and observed data (Raissi et al., 2020). PINNs have demonstrated success across many fields such as fluid mechanics (De Florio et al., 2021, 2022b; Cai et al., 2021), epidemiology (Schiassi et al., 2021a; Millevoi et al., 2024; Han et al., 2024; Kharazmi et al., 2021), cardiovascular mechanics (Sahli Costabal

70   et al., 2020; Yin et al., 2021; De Florio et al., 2025), and beyond (De Florio et al., 2024; Ahmadi Daryakenari et al., 2024), while extensions incorporating variational inference, dropout, and deep ensembles have enabled principled quantification of aleatoric and epistemic uncertainty under noisy or incomplete observations (Hüllermeier and Waegeman, 2021; Zou and Karniadakis, 2025; De Florio et al., 2025).

Despite their flexibility, PINN-based methods incur significant computational cost due to backpropagation through complex

75   loss terms and often require careful tuning of weighting coefficients between physics and data objectives. To alleviate these challenges, the Theory of Functional Connections (TFC) (Mortari, 2017b) provides an alternative "constrained expression" for differential-equation solutions that analytically satisfy boundary or initial conditions (Mortari, 2017a). The resulting con-

strained approximation permits rapid, training-free inference via direct least-squares solves of residuals at collocation points (Schiassi et al., 2022). Building on TFC, the X-TFC (Schiassi et al., 2021b) employs a single-layer random-projection network whose input weights and biases are randomly sampled once and held fixed. This technique has been widely used recently with promising performance in terms of accuracy and computational efficiency in several PIML frameworks (Fabiani et al., 2021; Fabiani, 2025; Fabiani et al., 2024; Dong and Li, 2021; Wang and Dong, 2024; Dwivedi and Srinivasan, 2020; Galaris et al., 2021).

X-TFC has been applied to both forward and inverse differential problems, demonstrating fast convergence and robustness in scenarios with limited, gappy, or noisy data (Osorio et al., 2025). Recent studies extend X-TFC to grey-box and black-box system identification (De Florio et al., 2024; Ahmadi Daryakenari et al., 2024) and uncertainty quantification (UQ) in physiological models, showing its capacity to decompose total uncertainty into aleatoric, epistemic, and model-form contributions (De Florio et al., 2025).

In this work, we review the conceptual foundations of PIML and X-TFC, apply X-TFC to an existing 1,475-hour vibration-based health-index (HI) dataset from a crack of the inner ring of a high-speed bearing in a 2.2-megawatt wind turbine gearbox (Bechhoefer and Dubé, 2020), and compare its effectiveness to classical physics-based fracture mechanics models. Finally, we show how our unified framework successfully also predicts RUL of bearings on different gearbox models using a HI generated from Supervisory Control and Data Acquisition (SCADA) measurements (Eftekhari Milani et al., 2025), broadening its applicability to a wider range of industrial assets. Once the effectiveness of X-TFC has been assessed, future work will explore the application of X-TFC to other bearing cracking problems in wind turbines, such as pitch bearing ring cracks shown in Figure 1(c). We also employ Monte Carlo (MC) X-TFC to quantify epistemic uncertainty for different scenarios of dataset sizes. The paper is structured as follows: Section 2 describes the vibration-based HI dataset from the 2.2-megawatt wind turbine gearbox, Section 3 describes classical physics-based fracture mechanics models, Section 4 describes PIML and the X-TFC method, Section 5 assesses the effectiveness of X-TFC combined with a classical physics-based fracture mechanics model to the vibration-based HI dataset, with UQ analysis, and to two similar cracks with HIs derived from only SCADA data, and Section 6 summarizes the analysis.

2 Crack Growth and Health Indicator Dataset

As described in Section 1, the dataset used in this paper is derived from an existing 1,475-hour (61-day) HI dataset from a crack of the inner ring of a high-speed bearing in a 2.2-megawatt wind turbine gearbox (Bechhoefer and Dubé, 2020). The HI represents a fusion of vibration-based condition indicators collected at irregular intervals (approximately every 10 minutes on average, with occasional 20–30 minute and multi-hour gaps), and used to extract features correlated with bearing damage. The HI was constructed such that a value of 0 represents a pristine bearing and a value of 1 indicates that it is appropriate to do maintenance. The RUL is then the time from the current HI until the HI is 1. Note that this nomenclature does not define a probability of failure for the component, nor does it mean that the component fails when the HI is 1. Instead, it suggests performing maintenance before collateral damage or cascading faults occur. By performing maintenance on a bearing prior to

shedding extensive material, costly gearbox replacement can be avoided, and the reliability of the gearbox can be restored to its design requirement. (Bechhoefer and Dubé, 2020).

Figure 2 (a) shows the HI calculated over the last 1,475 hours of operation of the wind turbine prior to the discovery of the bearing crack. Initially, the HI is very low; however, at 1,000 hours remaining, the HI began to rise, which corresponds to high

115 loads from a winter storm (Bechhoefer and Dubé, 2020). By the time 400 hours remaining is reached, the HI reached 0.5 on average with maximum values greater than 1. At 0 hours, the average HI exceeded 1, which triggered a borescope inspection of the bearing that uncovered the crack in the inner ring shown in Figure 2 (b). It is not known whether this crack originated from WECs; however, it has the same axial orientation on the inner ring and occurred in a high-speed bearing like many WECs in wind turbine gearboxes. This HI dataset will be used in the remainder of the paper to assess the effectiveness of X-TFC combined with a classical physics-based crack growth model.
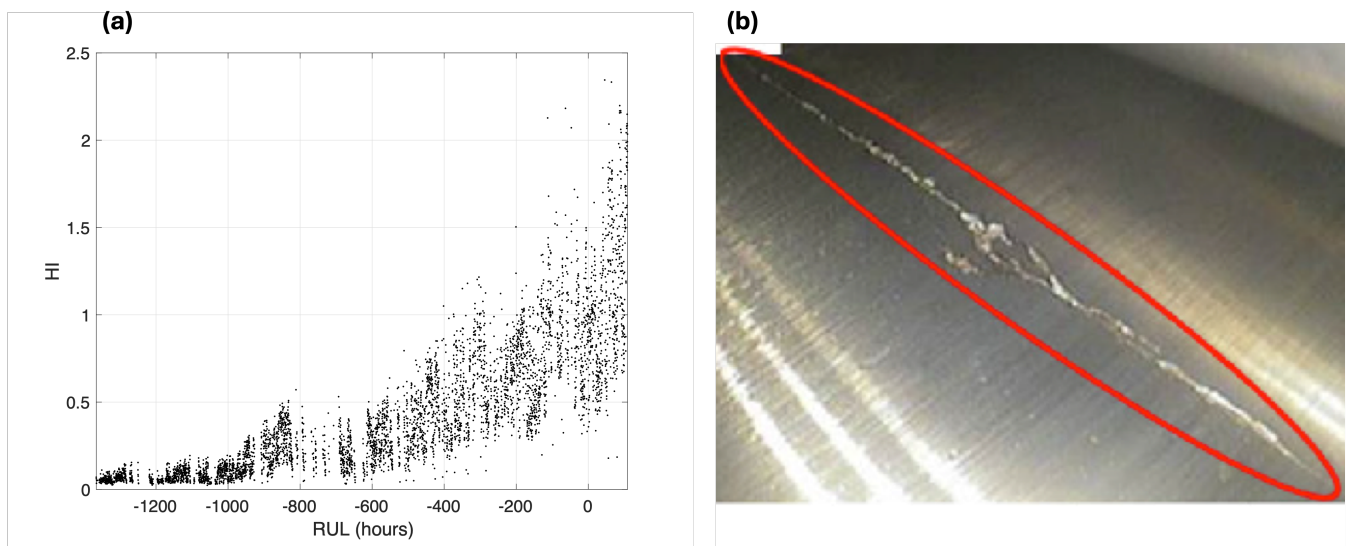
Figure 2. Example vibration-based HI dataset and bearing crack. (a) Calculated vibration-based HI (Bechhoefer and Dubé, 2020). (b) Crack in 2.2-megawatt wind turbine gearbox bearing (Bechhoefer and Dubé, 2020)

120

3 Fracture Mechanics Models

As stated in (Bechhoefer and Dubé, 2020), fatigue crack propagation is commonly described in terms of three fundamental displacement modes (Beer et al., 1999; Frost et al., 1999): Mode I (opening), Mode II (in-plane sliding), and Mode III (out-of-plane shear). In Mode I, the crack faces separate under tensile loading; in Mode II, the faces slide laterally within the plane

125 of the crack; and in Mode III, the faces move parallel to the crack front. By superposition of these modes, most general crack-tip displacement fields can be represented. Under linear-elastic assumptions, the crack surfaces are treated as traction-free boundaries so that external loads affect only the intensity of the stress field near the crack tip, which is quantified by the stress

intensity factor K. The factor K depends on the component geometry and loading; for simple geometries, it scales linearly with the applied nominal stress σ and with the square root of the crack length a (Beer et al., 1999). For example, the Mode-I

130   stress intensity factor for a crack of length a can be expressed as

$$K = \sigma\sqrt{\pi a}\,\alpha, \tag{1}$$

where σ is the far-field (gross) stress and α is a dimensionless shape factor. When K is known, stresses and strains near the crack tip can be calculated. For instance, the normal stress in the x-direction at a point (r, θ) measured from the crack tip is (see Frost et al. (1999))

135   $$\sigma_x(r, \theta) = \frac{K}{\sqrt{2\pi r}} \cos\left(\frac{\theta}{2}\right) \left[1 - \sin\left(\frac{\theta}{2}\right) \sin\left(\frac{3\theta}{2}\right) \right] \tag{2}$$

which illustrates the classical inverse-square-root singularity of the stress field. This formulation of the crack-tip field in terms of a single parameter K is a cornerstone of linear elastic fracture mechanics. In fatigue crack growth analysis, the driving force is taken as the stress intensity range ΔK (the difference between maximum and minimum K in a load cycle). Experimental evidence (Frost et al., 1999) shows that ΔK, not the maximum K, governs crack advance, and if ΔK is held constant, then

140   the crack growth rate becomes steady.

3.1   Linear Elastic Fracture Mechanics

For many engineering materials (e.g., steels in gears and bearings) subjected to cyclic tensile loading, the incremental crack growth per cycle is Mode I, and it follows a power-law relation (i.e. Paris' law (Pugno et al., 2006)). A commonly used form is

145   $$\frac{\mathrm{d}a}{\mathrm{d}N} = D(\Delta K)^m, \tag{3}$$

where $\mathrm{d}a/\mathrm{d}N$ is the crack extension per load cycle, D is a material-specific constant, and m is the crack-growth exponent (typically $3 \le m \le 5$). Substituting the expression $\Delta K = 2\sigma\sqrt{\pi a}\alpha$ into Eq. (3) yields

$$\frac{\mathrm{d}a}{\mathrm{d}N} = D\left[2\sigma\sqrt{\pi}\alpha \right]^m a^{m/2}. \tag{4}$$

Rearranging Eq. (4) and integrating both sides gives the total number of cycles N required to grow the crack from an initial

150   length a_0 to some length a_f

$$N = \int_{a_0}^{a_\mathrm{f}} \frac{a^{-m/2}}{D\left(2\sigma\sqrt{\pi}\alpha \right)^m} \, \mathrm{d}a \tag{5}$$

Performing the integration yields

$$N = -\frac{\mathrm{d}N}{\mathrm{d}a} a_0 - \frac{a_\mathrm{f}(a_0/a_\mathrm{f})^{m/2}}{m/2 - 1} \tag{6}$$

provided $m \neq 2$. In particular, if one assumes $m = 4$, Eq. (6) reduces to

$$\frac{\mathrm{d}N}{\mathrm{d}a} = -\frac{N}{a_0 \ln\left(\frac{a_\mathrm{f}}{a_0}\right)} \tag{7}$$

Under steady, cyclic loading (e.g., constant-speed operations for helicopter gearboxes or wind turbines in region 3), the number of cycles N is directly proportional to operating time. If the prognostic HI is assumed proportional to the accumulated damage (so that advancing crack length corresponds to decreasing HI), then Eq. (7) provides a RUL estimate by relating crack growth to the elapsed time or number of cycles. We can select $a_\mathrm{f} = 1$, defining the need to perform maintenance when the HI is 1.

3.2 Head's Theory

Head's theory proposes a simplified model of the crack-tip stress field in which the material ahead of the crack tip is treated as an assembly of parallel elastic bars. Each bar of modulus E carries the applied remote stress σ both directly and through shear transfer. The resulting crack-growth law can be cast in a form similar to Eq. (4), but with a higher effective exponent

$$\frac{\mathrm{d}a}{\mathrm{d}N} = \frac{2E\sigma^3 a^{3/2}}{3E_\mathrm{W}(\Delta\sigma)D} \propto K^3 \tag{8}$$

In effect, Head's model predicts a crack advance that corresponds to $m = 6$ in the Paris' law form. Although the intermediate steps involve material and geometric parameters, the final integrated form resembles Eq. (6) with $m = 6$. In particular, inversion and integration analogous to the linear-elastic case yield an alternate RUL expression

$$\frac{\mathrm{d}N}{\mathrm{d}a} = \frac{N}{a_0 \left(2a_\mathrm{f} - \frac{2}{\sqrt{a_0}}\right)} \tag{9}$$

3.3 Dislocation Theory

In the case of Mode-III (anti-plane shear) dominant loading, dislocation (or plastic dislocation) theory may apply. Here the crack-tip plastic zone is modeled as a continuous distribution of dislocations. Crack advance is assumed to occur when the integrated plastic slip reaches a threshold. The associated growth law can be expressed as

$$\frac{\mathrm{d}a}{\mathrm{d}N} = \frac{a^2 \sigma_\mathrm{max}^4}{DE\sigma^3} \tag{10}$$

This equation implies that crack advance per cycle grows with decreasing crack length, similar in form to the Paris' law in Eq. (3) with an effective exponent $m = 2$. By inverting, integrating, and changing terms we get

$$\frac{\mathrm{d}N}{\mathrm{d}a} = -\frac{N}{a_0(2a_\mathrm{f} - 2\sqrt{a_0})} \tag{11}$$

3.4 Comparison of Crack Growth Models

Figure 3 compares the three models described in the previous sections when applied to the full 1,475 hour vibration-based HI dataset described in Section 2. For this comparison, the values of a_0 and a_f are 0.05 and 1, respectively. With the given crack

growth exponent m assumed in each model, Head's theory best matches the slow rise of the calculated HI data, especially over the last 700 hours. In comparison, the linear elastic and dislocation theory models have a very slow rise, followed by a much steeper rise thereafter that would provide very little warning of the necessity of an impending maintenance action. Thus, in the remainder of this paper, we choose to use the Head's theory model in combination with X-TFC.
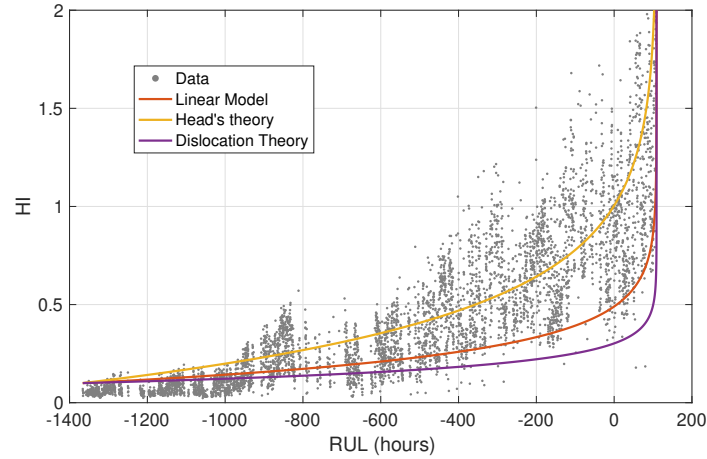
Figure 3. Solutions of the differential equations given by linear, Head's theory, and Dislocation theory models compared to the calculated vibration-based HI.

4 Physics-Informed Machine Learning

Physics-Informed Machine Learning (PIML) is an emerging framework that integrates known physical laws into the training of data-driven models. By embedding governing differential equations, boundary conditions, and other domain-specific constraints directly into the loss function, PIML enables models to learn from both observational data and prior physical knowledge. This approach contrasts with purely data-driven methods, which may struggle when observations are scarce, noisy, or incomplete.

Consider a physical system governed by partial differential equations,

$$\mathcal{N}[u(\mathbf{x}), \mathbf{x}] = 0, \quad \mathbf{x} \in \Omega, \tag{12}$$

where $u(\mathbf{x})$ is the state variable and $\mathcal{N}[\cdot]$ is a differential operator defined on the domain Ω. Traditional data-driven models seek a surrogate $\hat{u}(\mathbf{x}; \theta)$ parameterized by θ that minimizes a data loss

$$\mathcal{L}_{\text{data}}(\theta) = \frac{1}{N_{\text{d}}} \sum_{i=1}^{N_{\text{d}}} \left| \hat{u}(\mathbf{x}_i; \theta) - u_i^{\text{obs}} \right|^2, \tag{13}$$

where $\{(\mathbf{x}_i, u_i^{\text{obs}})\}_{i=1}^{N_{\text{d}}}$ are N_{d} observed data points. In contrast, PIML augments this objective with a physics-based penalty

$$\mathcal{L}_{\text{physics}}(\theta) = \frac{1}{N_{\text{p}}} \sum_{j=1}^{N_{\text{p}}} \left| \mathcal{N}[\hat{u}(\mathbf{x}_j; \theta), \mathbf{x}_j] \right|^2, \tag{14}$$

where $\{\mathbf{x}_j\}_{j=1}^{N_{\text{p}}}$ are collocation points chosen in Ω. The physical loss is made by the residual of the differential equation, which needs to be minimized. In our case, it is the residual of the Eq. (9):

$$\frac{\mathrm{d}N}{\mathrm{d}a} - \frac{N}{a_0 \left(2a_{\text{f}} - \frac{2}{\sqrt{a_0}} \right)} = 0 \tag{15}$$

The total loss is then given by

$$\mathcal{L}(\theta) = \lambda_{\text{data}} \cdot \mathcal{L}_{\text{data}}(\theta) + \lambda_{\text{physics}} \cdot \mathcal{L}_{\text{physics}}(\theta), \tag{16}$$

with weighting coefficients λ_{data} and λ_{physics} balancing data fidelity against physical consistency.

By enforcing Eq. (12) during training, PIML offers several advantages when dealing with realistic datasets:

– **Regularization via Physics:** The physics-based term $\mathcal{L}_{\text{physics}}$ acts as a principled regularizer, preventing overfitting to noise in u_i^{obs}.

– **Data Efficiency:** Incorporating prior knowledge reduces the reliance on large volumes of labeled data, enabling accurate learning even when N_{d} is small or spatial coverage is incomplete.

– **Robustness to Missing Observations:** The model can be evaluated at arbitrary collocation points, filling gaps in the observational domain and providing continuous field reconstructions.

– **UQ:** Extensions of PIML integrate Bayesian or ensemble approaches to quantify epistemic and aleatoric uncertainties, further improving reliability under data scarcity.

4.1 eXtreme Theory of Functional Connections (X-TFC)

The PIML framework used in this work is based on the X-TFC algorithm (Schiassi et al., 2021b), which differs from classic PINN for two main feature introduced: the TFC (Mortari, 2017b, a; Mortari et al., 2019), which allows to satisfy initial and boundary conditions of the differential equations analytically, and random-projection neural networks (Fabiani, 2025) trained via Extreme Learning Machine algorithm (Huang et al., 2006), avoiding thus the need of the computationally expensive back-propagation algorithm. The TFC constructs an explicit approximation of the solution of a differential equation that identically satisfies given boundary or initial conditions. For a general initial value problem of the form

$$\frac{\mathrm{d}x}{\mathrm{d}t} = f\big(x(t), t\big), \quad x(0) = x_0, \tag{17}$$

TFC represents the solution of the differential equation via a so-called 'constrained expression'

$$x(t; \beta) = g(t; \beta) - g(0; \beta) + x_0, \tag{18}$$

where $g(t;\beta)$ is any sufficiently smooth function parameterized by coefficients β. In the X-TFC framework (De Florio et al., 2022a), $g(t;\beta)$ is chosen as a random-projection single-layer neural network whose input weights and biases $\{w_j, b_j\}_{j=1}^{L}$ are randomly sampled once and held fixed. Specifically,

$$225 \quad g(t;\beta) = \sum_{j=1}^{L} \beta_j \sigma\left(w_j, t + b_j\right) = \begin{bmatrix} \sigma(w_1 t + b_1) \\ \sigma(w_2 t + b_2) \\ \vdots \\ \sigma(w_L t + b_L) \end{bmatrix}^{T}, \tag{19}$$

and

$$\sigma^T(0) = \begin{bmatrix} \sigma(b_1) \\ \sigma(b_2) \\ \vdots \\ \sigma(b_L) \end{bmatrix} = \sigma_0^T, \tag{20}$$

where $\sigma(\dots)$ is the neural network's activation function, and β is the vector of the output weights of the neural network to be learned.

230   Thus, a parametric approximation of the solution of the ordinary differential equation and its derivative in time are

$$x(t;\beta) = [\sigma - \sigma_0]^T \beta + x_0, \tag{21}$$

and

$$\dot{x}(t;\beta) = \dot{\sigma}(t)^T \beta, \tag{22}$$

A least-squares formulation minimizes the residuals of the ordinary differential equation at the collocation points $\{t_i\}$ and any
235   available data. Denoting the vector of ordinary differential equation residuals loss

$$\mathcal{L}_{\text{physics}} = \frac{\mathrm{d}x}{\mathrm{d}t} - f\left(x(t), t\right) \tag{23}$$

and data loss

$$\mathcal{L}_{\text{data}} = x(t;\beta) - \hat{x}(t;\beta) \tag{24}$$

(where $\hat{x}(t;\beta)$ represents the observation data), by $\mathcal{L} = \mathcal{L}_{\text{physics}} + \mathcal{L}_{\text{data}}$, the update step in a Gauss–Newton scheme to deter-
240   mine the k-th iterate of the unknown coefficients β is

$$\Delta \beta^k = -\left[\mathcal{J}^T \mathcal{J}\right]^{-1} \mathcal{J}^T \mathcal{L}, \tag{25}$$

where $\mathcal{J}$ is the Jacobian of losses with respect to coefficient vector β. A schematic of the X-TFC framework is shown in Figure 4. X-TFC offers rapid, training-free inference by exploiting random projections and a direct least-squares solver, making it well suited for online estimation and UQ.
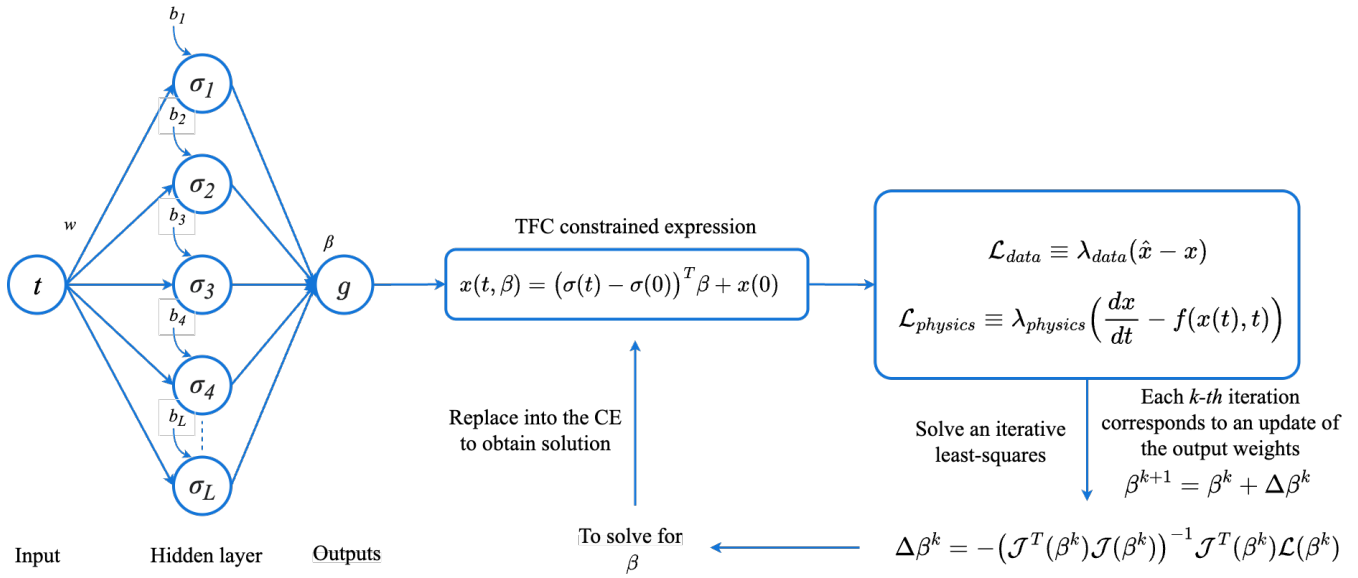
Figure 4. Schematic of the X-TFC algorithm that includes data and physics in the loss function. Input weights and biases are randomly selected. The last step solves the least-squares problem.

5 Results

As previously mentioned, the physics loss consists of the residual of the differential equation modeled by the Head's theory as a fracture mechanics model, where the values of a_0 and a_f are 0.05 and 1, respectively. The simulations are run with 1,000 collocation points of the neural network's input, sampled equally spaced between -1 and 1, and with five hidden neurons. The number of collocation points and neurons are user-selected to achieve a good compromise between accuracy and computational expense. The input weights and bias are sampled with a uniform distribution between -1 and 1, and the selected activation function is *tanh*. Figure 5 reports the results for all our simulations. Each plot represents a different data size scenario to understand the advantage of using PIML, especially when data availability is scarce. For each data size scenario, we performed the RUL prediction for different weights given to the differential equation, such as $\lambda_\mathrm{physics} = [0.0, 0.25, 0.50, 0.75, 1.0]$, where scenario $\lambda_\mathrm{physics} = 0$ is the case in which we do not have knowledge of physics, so we only perform data-driven regression.

The scatter data points in blue represent the full dataset we have for the bearing under study, while the red scatter data points are those used in our framework, simulating conditions of missing data. The time domain, which represents the RUL in hours, is calibrated to correspond to zero when the HI reaches a value of 1, which means that there is a high probability of bearing damage and need to perform maintenance, obtained with the scenario of full knowledge of physical ($\lambda_\mathrm{physics} = 1$) and full availability of data (see Figure 5(a)). This gives us our ground truth solution. The results in Figure 5(a) represent the case in which we assume to know the data throughout the entire domain, allowing us to see the effect of the weights on the physics

loss to the RUL prediction. When we assume no knowledge of physics, the data alone is sufficient to have a good prediction of RUL, which differs from our ground truth of about 50 hours, as reported in Table 1.

| Data Availability | RUL error (hours) | | | | | Time |
| | Physics Weight | | | | | |
(%)	1.0	0.75	0.50	0.25	0.0	(ms)
100	0.00	9.74	24.78	44.39	55.57	45
75	22.15	42.09	73.27	110.98	133.19	35
50	42.25	75.50	131.50	208.59	240.30	25
25	54.48	102.51	202.23	409.15	597.23	15

Table 1. Error values of RUL prediction in hours for different combinations of data availability and physics weights. On the right column, the computational time in milliseconds is reported.

The plots in Figures 5 (b),(c), and (d) show a linear trend of the deviation of the RUL prediction, as the availability of data and knowledge of physics decrease, as expected. This is clear evidence of the improvement brought by physics that reduces the search space previously based only on data. In fact, it can be observed that even with low data availability (25%), the RUL prediction has an error of only about 50 hours when full weight (physics weight = 1.0) is given to the loss of the differential equation. However, with no contributions from physics (physics weight = 0.0), the RUL predication error increases to ≈ 600 hours. In Figure 6, the RUL error surface is plotted, obtained with 9,191 simulations (91 dataset sizes from 0.1% to 100%, and 101 physics weights from 0 to 1), showing a mirror-like behavior that explains the equal importance of both data and physics, fundamental ingredients of PIML, and the advantage of using both of them in synergy to obtain the maximum performance gain. The simulations in this work were performed with a laptop MacBook Pro, Apple M3 Pro, 12 cores, 36GB, and MATLAB R2024b.

5.1 Uncertainty Quantification

To quantify epistemic uncertainty, we employed a simple yet scalable MC strategy, which is called MC X-TFC (De Florio et al., 2025). In this method, we independently train multiple X-TFC models on synthetic data corrupted by random noise drawn from a known distribution, and each model starts from a different random initialization of weights and biases. MC X-TFC closely parallels the "deep ensembles" frameworks, as proposed in Lakshminarayanan et al. (2017), which has proven highly effective for neural-network UQ, even when that uncertainty arises solely from random parameter initialization (Ganaie et al., 2022; Rahaman et al., 2021; Wenzel et al., 2020; Gawlikowski et al., 2023; Psaros et al., 2023; Zou et al., 2024b; Pickering et al., 2022; Zou et al., 2024a; Zhang et al., 2024). We define epistemic uncertainty as the variability in the predicted state vector that arises from two sources:

1. Model-parameter uncertainty, that is, the fluctuations in the coefficients β induced by the random sampling of network weights w and biases b in Eq. (19);
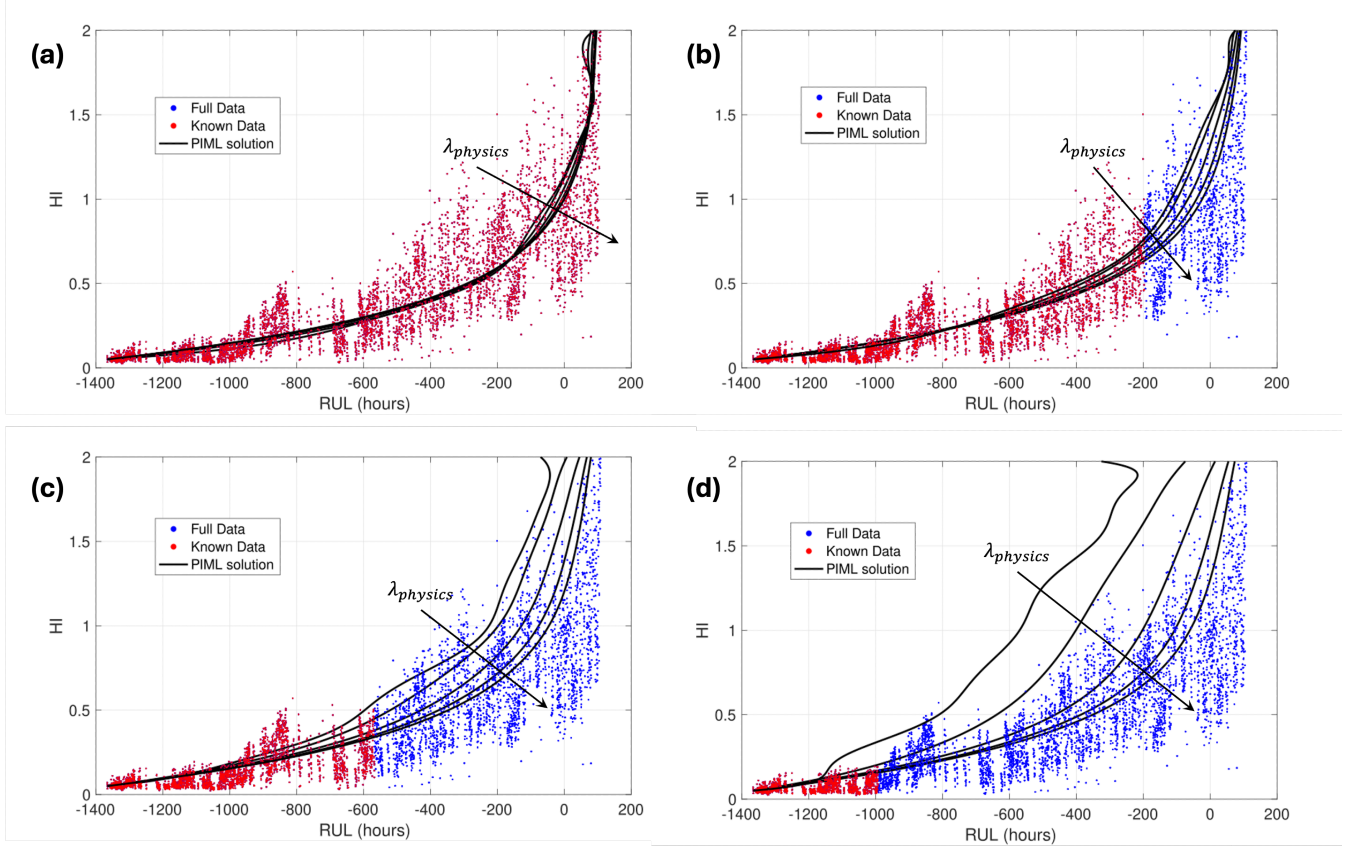
Figure 5. Results of PIML for different dataset sizes. For each scenario, different weights ($\lambda_{\text{physics}} = [0.0, 0.25, 0.50, 0.75, 1.0]$) on the equation loss are given to understand the improvement in the RUL forecasting brought by the physics knowledge. (a) Full dataset available: 100%. (b) Partial dataset available: 75%. (c) Partial dataset available: 50%. (d) Partial dataset available: 25%.

2. Data-driven uncertainty, that is, the variability in the observed data under our noise model, which we also propagate into the initial condition a_0 of the HI, by perturbing it with zero-mean Gaussian noise, with standard deviations of 20%.

In our analysis, we evaluate the predictive performance of the model using the Mean of Error (ME), the Standard Deviation of Error (SDE), and Confidence Interval (CI) of signed errors, where errors retain their direction (positive or negative). All error values and derived statistics are expressed in hours, and the simulations are performed with 100 MC realizations and physical loss weight equal to 1. The ME is calculated as

$$\text{ME} = \frac{1}{m_{\text{c}}} \sum_{i=1}^{m_{\text{c}}} e_i, \tag{26}$$

where e_i (in hours) is the error for the i-th sample (predicted minus true value), and m_{c} is the total number of samples. The ME indicates systematic bias: a positive ME means earlier predictions of failure, a negative ME means later predictions of failure.
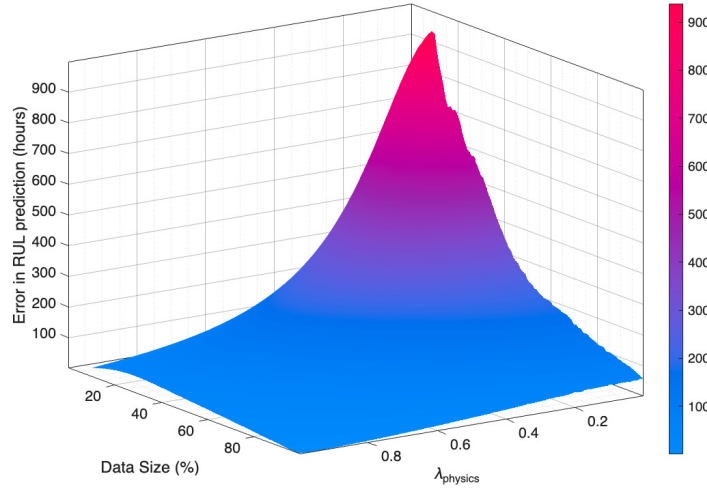
Figure 6. Error in RUL prediction (hours) for different scenarios of dataset size and weights on physics loss.

The SDE is given by:

$$\text{SDE} = \sqrt{\frac{1}{m_{\text{c}}-1}\sum_{i=1}^{m_{\text{c}}}(e_i - \text{ME})^2}, \tag{27}$$

295 which measures the spread of signed errors around the ME. A larger SDE (in hours) indicates greater variability in the prediction deviations. We also compute non-parametric percentiles of the e_i distribution to capture coverage analogous to Gaussian confidence levels:

- **68% Interval ("1σ"):** 16th to 84th percentiles of e_i.

- **95% Interval ("2σ"):** 2.5th to 97.5th percentiles of e_i.

300 - **99.7% Interval ("3σ"):** 0.15th to 99.85th percentiles of e_i.

These intervals indicate where the central 68%, 95%, and 99.7% of signed errors lie, respectively, preserving information on whether prediction errors tend to be positive or negative. All these error metrics for epistemic UQ are reported in Table 2, for different dataset sizes. As expected, the absolute values of ME and SDE increase with the decrease in dataset size. The negative values of ME represent an error in an overly early prediction of failure. Computational times also decrease with the decrease of

305 dataset size, because of the smaller number of points in the data loss function. The CIs reported in Table 2 are also qualitatively shown in Figure 7 for the three confidence percentages. Despite the inherent stochasticity of random model initialization and noisy observations, the resulting CIs remain moderate (e.g., ±10 hours for a 68% CI on the full dataset). This behavior stems primarily from the strong physics constraint given by the Head's theory differential equation, which tightly regularizes the

| Data Availability (%) | RUL error (hours) | | CI of error (hours) | | | Time (s) |
	ME	SDE	68%	95%	99.7%	
100	7.43	11.28	[-3.16, 17.52]	[-9.92, 34.45]	[-13.22, 38.86]	7.4
75	-18.09	14.63	[-31.80, -4.87]	[-40.76, 16.67]	[-45.13, 22.22]	5.7
50	-47.72	19.42	[-66.03, -30.05]	[-78.19, -2.08]	[-84.28, 4.97]	4.9
25	-73.54	24.87	[-97.05, -50.68]	[-113.20, -16.01]	[-121.35, -7.46]	3.2

Table 2. Error values of RUL prediction in hours for different combinations of data availability while quantifying epistemic uncertainty. The error metrics in hours are the Mean of Error (ME), the Standard Deviation of Error (SDE), and the Confidence Interval (CI) of error. Last column reports the computational time in seconds for 100 MC realizations.
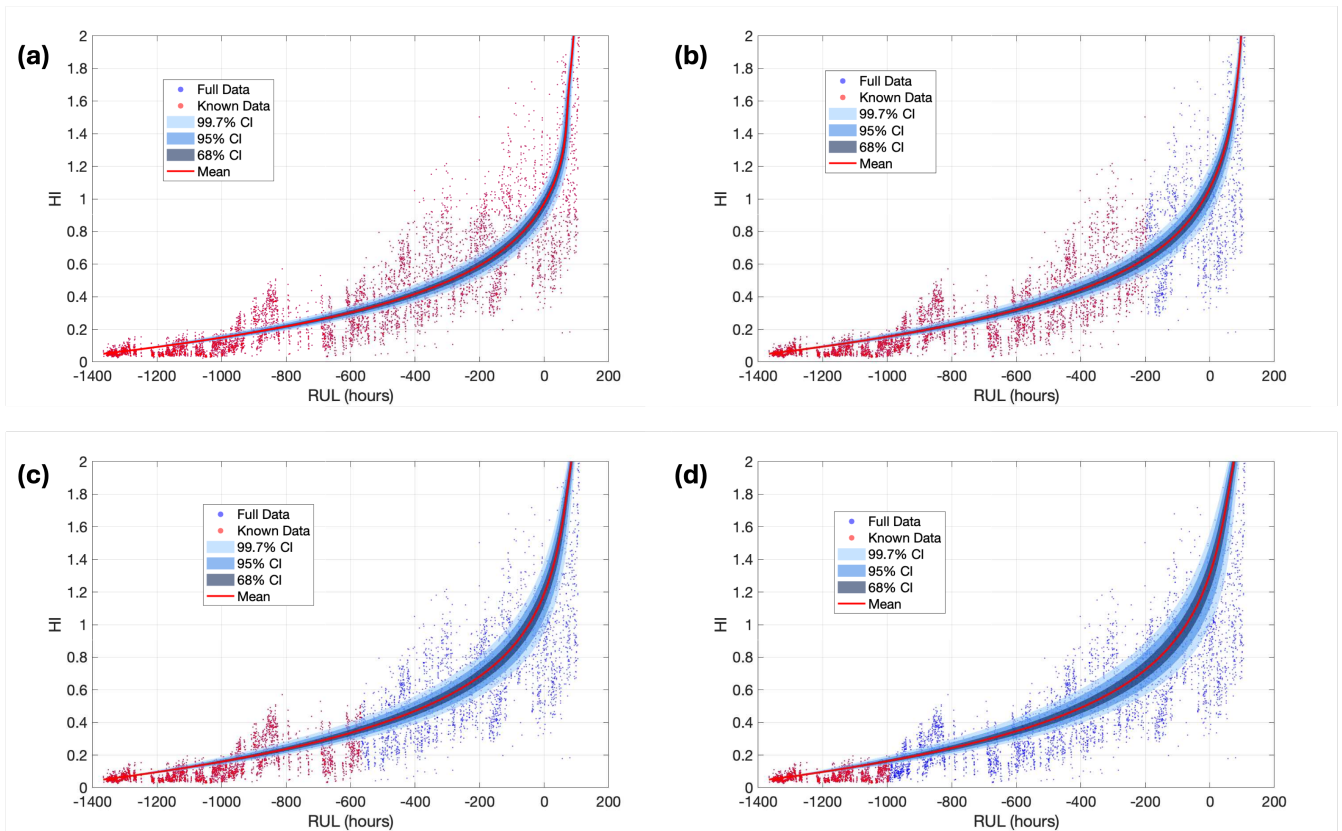
Figure 7. Epistemic quantification for different dataset sizes, considering $\lambda_{\text{physics}} = 1$ on the equation loss, and three CIs with 67%, 95%, and 99.7% confidence obtained with 100 MC realizations. (a) Full dataset available: 100%. (b) Partial dataset available: 75%. (c) Partial dataset available: 50%. (d) Partial dataset available: 25%.

solution and limits divergence across ensemble members. As such, once model-induced variability is largely suppressed by
310 physics-informed regularization, the spread in predictions stays relatively small, highlighting the effectiveness of embedding
physical laws into the learning process.

Since the epistemic uncertainty is reducible, the CIs presented in Figure 7 are obtained with optimized parameters that reduce
the epistemic uncertainty as low as possible. This reduction depends on the many hyperparameters selected, such as activation
function, number of collocation points, number of neurons, least-squares iteration tolerance, and random sampling ranges. An
315 ablation study on the behavior of epistemic uncertainty depending on the choice of the random initialization of the neural
network input weights and bias is shown in Figure 8. The three plots show that the different choices of input weights and bias
ranges can lead to wider CIs for the RUL prediction, and by choosing the optimal hyperparameter set up, the epistemic (model)
uncertainty can be reduced. Further reduction in prediction spread would therefore require enhancements in data quality, such
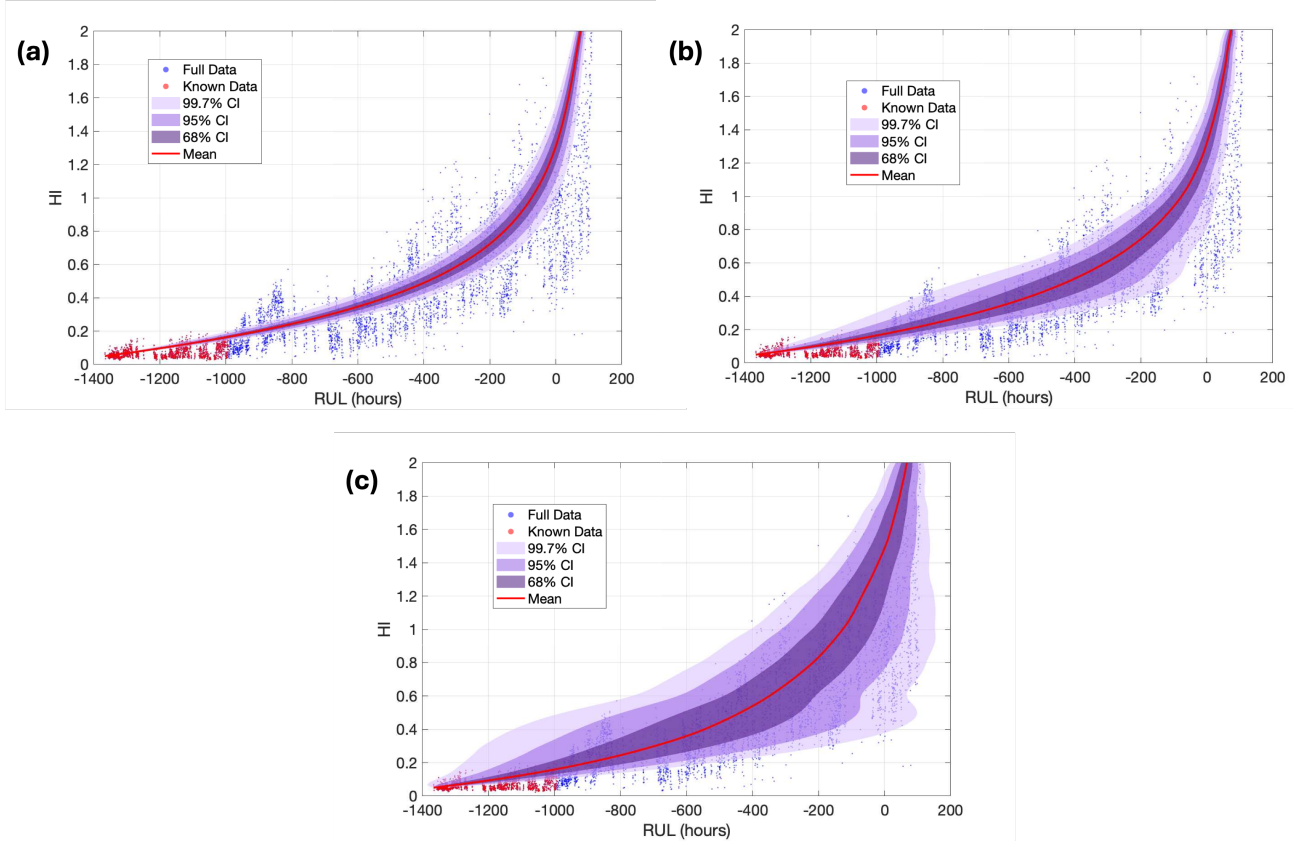
Figure 8. Ablation study on epistemic quantification for different random initialization of input weights and bias ranges, considering $\lambda_{\text{physics}} =$ 1 on the equation loss, partial data size (25%), and three CIs with 67%, 95%, and 99.7% confidence obtained with 100 MC realizations. Random initialization of input weights and bias in range (a) [-1 ,1], (b) [-10 ,10],(c) [-20 ,20].

as using higher resolution measurements, to minimize noise-driven variability. Moreover, the aleatoric uncertainty is inherently

320 irreducible (Hüllermeier and Waegeman, 2021) and, in this case, is particularly high due to the very noisy data. In practical terms, this means that there is a minimum level of predictive uncertainty, driven by data noise, that cannot be further reduced by improving the model. Consequently, even a perfect model would still exhibit this baseline uncertainty dictated by the natural variability of the observations. This is shown in Figure 9, where the CI of the aleatoric uncertainty is plotted against the noisy dataset. Because the dataset exhibits heteroscedastic noise, that is, noise amplitude varies over time, we first estimate the global

325 noise scale by analyzing the full HI series. We apply a LOESS (Locally Estimated Scatterplot Smoothing) method to find a smoothed trend of the HI over time and then compute the residuals using the noisy dataset to capture fluctuations. Then, the overall noise level is approximated as the standard deviation of the residuals. Finally, the aleatoric uncertainty band in Figure 9 is generated by plotting the smoothed trend $\pm 2\sigma$, reflecting the heteroscedastic, irreducible noise-driven uncertainty in the HI measurements.
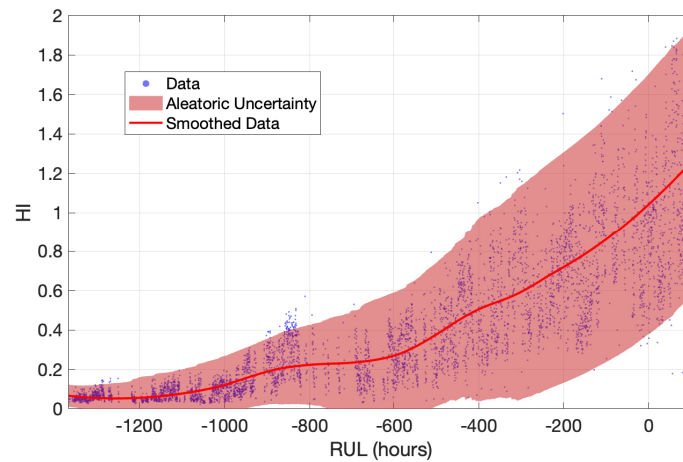
Figure 9. Aleatoric uncertainty band estimated from the vibration-based HI. The shaded region represents $\pm 2\sigma$ bounds based on the time-varying standard deviation of the residuals between noisy data and smoothed data. This illustrates the heteroscedastic, irreducible uncertainty arising from measurement noise.

330 ## 5.2 Application to SCADA-Based Health Index

To demonstrate the generality and adaptability of the X-TFC framework, we further apply our methodology to a derived HI from SCADA data described in Eftekhari Milani et al. (2025). The SCADA dataset is from two 1.5-megawatt wind turbines, with measurements available in a 10-minute time frame. These turbines have identical gearbox configurations, with one experiencing an intermediate-speed shaft bearing axial crack and the other experiencing a high-speed shaft bearing axial crack. These

335 failures are the first gearbox-related failures in these wind turbines, enabling an analysis of the degradation from a pristine state to failure.

In the data pre-processing step, the non-operational data points with a zero rotor speed or a negative power are removed. In addition, the time stamps with missing or erroneous values due to sensor problems are dropped. Then, a daily averaging is performed, decreasing the missing data percentage from 14% in the 10-minute time frame to around 2% in the daily time
340 frame. This averaging increases data continuity and reduces the computational burden.

To obtain the HIs, the approach proposed in Eftekhari Milani et al. (2025) is used. First, for each turbine, the gearbox high-speed shaft bearing temperature in the healthy state is modelled by training a convolutional neural network using the remaining six SCADA channels during the first year of operation, in which the turbines are expected to be healthy. Then, the modelled healthy-state gearbox bearing temperature values are subtracted from the recorded values, obtaining the residuals. They are
345 then normalised according to the one-year training period to obtain the residual z-scores. These z-scores, along with the four operational and environmental signals of power, rotor speed, wind speed, and ambient temperature, are input to a convolutional autoencoder (CAE)-based unsupervised HI construction model to build the HIs. A leave-one-out cross-validation is performed, where one wind turbine is set aside for testing and the remaining five wind turbines are used for training the CAE. The trained CAE is then used to obtain the test wind turbine's HI.

350 While purely data-driven methods like LSTM and XGBoost showed promise for SCADA data, their precision was limited, often resulting in a high number of false alarms (Desai et al., 2020). To address this limitation, we applied X-TFC to SCADA-based HIs of two turbines (Eftekhari Milani et al., 2025), embedding the same Head's theory differential equation used previously for vibration-based data. The physics model acts as a regularizer, enforcing consistency between the SCADA-derived HI trajectory and the expected fatigue crack growth behavior. In Figure 10, epistemic uncertainty of the X-TFC RUL
355 prediction for the SCADA data-based HI is reported. When trained on only 25% of the available SCADA HI trajectory, X-TFC obtains failure predictions with a confidence of about 10 days. Importantly, the physics-informed structure of X-TFC helped to suppress model overfitting, compared to traditional machine learning models applied directly to SCADA signals (Desai et al., 2020). This confirms the versatility and applicability of the X-TFC framework to both vibration- and SCADA data-based HIs. This widens the scope of prognostics applications for wind turbines, enabling early and reliable failure prediction even when
360 only typical SCADA data are accessible.

6 Conclusions

In this work, we have introduced and demonstrated the application of the X-TFC as a PIML framework for predicting the RUL of wind turbine gearbox bearing cracks. By embedding Head's fracture-mechanics model directly into the learning objective, X-TFC balances fidelity to HIs based on both vibration and SCADA data with consistency to the differential crack-growth
365 law. In particular, we calibrated the baseline PIML time-domain solution to be zero for the full 1,475-hour vibration-based HI dataset from a cracked high-speed bearing in a 2.2-megawatt wind turbine gearbox and full physics weighting ($\lambda_{\text{physics}} = 1$) to a value of HI equal to 1, indicating a high possibility of bearing damage and need to perform maintenance. Lower values of weights and purely data-driven regression ($\lambda_{\text{physics}} = 0$) still lead to a good accuracy in RUL prediction with less than an error of 60 hours. Even when only 25% (or less) of the vibration-based HI data are available, heavy weighting on the physics loss
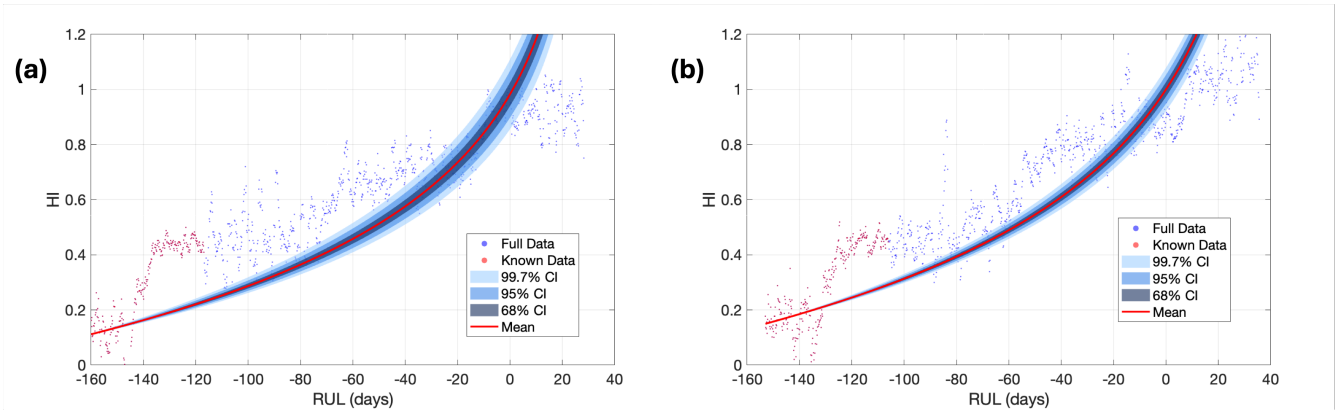
Figure 10. Epistemic uncertainty quantification with SCADA-based HI of two different gearbox bearings, considering $\lambda_{\text{physics}} = 1$ on the equation loss, and three CI with 67%, 95%, and 99.7% confidence obtained with 100 MC realizations. (a) Gearbox bearing #1. (b) Gearbox bearing #2.

370   yields RUL errors below 60 hours, a four-fold improvement over the data-only case (≈ 600 hours). This demonstrates X-TFC's ability to leverage domain knowledge to fill gaps in noisy, sparse, or intermittent monitoring data, in very low computational times (in the order of milliseconds). To further enhance the interpretability and robustness of these predictions, we introduced a MC ensembling strategy (MC X-TFC) to quantify epistemic uncertainty. This approach produces CIs for the estimated RUL by propagating uncertainty due to both model initialization and measurement noise, thereby enabling probabilistic forecasting.

375   These intervals provide insight into the reliability of each prediction and help operators assess risk in real-time maintenance planning. Finally, as shown for this case, the irreducible aleatoric uncertainty is high due to the natural variability of the HI data.

    These results demonstrate the potential of X-TFC as a powerful, data-efficient, and computationally light methodology for bearing fatigue-crack prognostics in rotating machinery. By guaranteeing physically plausible degradation trajectories regard-

380 less of data sparsity, X-TFC enables earlier and more reliable maintenance decisions, with the potential to reduce unplanned downtime and O&M costs for wind turbine gearboxes. Our study, in its early stage, focuses on HI data derived from bearing cracks and employs Head's theory as the governing crack-growth law. By validating our methodology using HI data generated from both vibration and SCADA datasets, we confirm its versatility to bearings of multiple gearbox manufacturers, models and ratings. To capture more realistic wind-turbine bearing behavior, our PIML framework should be validated against larger,

385 more diverse datasets. These models would enable real-time monitoring of bearing health by continuously ingesting streaming data and updating predictions as new information arrives.

    Future extensions should also examine fatigue cracking in pitch bearing rings, which may require more comprehensive fracture-mechanics models such as multi-phase Paris laws or dislocation-theory blends, to capture more complex and realistic nonlinear crack–load interactions. Continued development of UQ capabilities, such as MC X-TFC, will also be critical for

390 producing probabilistic RUL estimates with actionable confidence bounds, especially in high-stakes decision environments.

Author contributions. Writing – original draft preparation, formal analysis, software, data curation, visualization by MD; writing – review and editing by MD, GA, JK, AEM, DZ, and SS; data acquisition by AEM and DZ; conceptualization by MD, JK, and SS; funding acquisition and project administration by SS.

Competing interests. At least one of the (co-)authors is a member of the editorial board of Wind Energy Science.

395

Code availability. The codes used for the analyses and experiments presented in this study are available from the authors upon request.

Acknowledgements. The authors would like to thank Dr. Eric Bechhoefer for providing the vibration-based HI dataset used in this research.

References

Ahmadi Daryakenari, N., De Florio, M., Shukla, K., and Karniadakis, G. E.: AI-Aristotle: A physics-informed framework for systems biology gray-box identification, PLOS Computational Biology, 20, e1011 916, https://doi.org/10.1371/journal.pcbi.1011916, 2024.

Bechhoefer, E. and Dubé, M.: Contending Remaining Useful Life Algorithms, in: Annual Conference of the PHM Society, vol. 12, pp. 9–9, 2020.

Bechhoefer, E., Xiao, L., and Zhang, X.: Remaining Useful Life Calculation of a Component using Hybrid Fatigue Crack Model, in: Annual Conference of the PHM Society, vol. 13, https://doi.org/10.36001/phmconf.2021.v13i1.3062, 2021.

Beer, F., Johnston, E., and DeWolf, J.: Mechanics of materials, 5th SI Edition, Stress, 1, 1–12, 1999.

Cai, S., Mao, Z., Wang, Z., Yin, M., and Karniadakis, G. E.: Physics-informed neural networks (PINNs) for fluid mechanics: A review, Acta Mechanica Sinica, 37, 1727–1738, https://doi.org/10.1007/s10409-021-01148-1, 2021.

Clark, C., Guo, Y., Sheng, S., and Keller, J.: Effects of Bearing Clearance, Oil Viscosity, and Temperature on Bearing White-Etching Cracks, NREL/TP-5000-85917, Golden, CO: National Renewable Energy Laboratory, 2023.

De Florio, M., Schiassi, E., Ganapol, B. D., and Furfaro, R.: Physics-informed neural networks for rarefied-gas dynamics: Thermal creep flow in the Bhatnagar–Gross–Krook approximation, Physics of Fluids, 33, https://doi.org/10.1063/5.0046181, 2021.

De Florio, M., Schiassi, E., and Furfaro, R.: Physics-informed neural networks and functional interpolation for stiff chemical kinetics, Chaos: An Interdisciplinary Journal of Nonlinear Science, 32, https://doi.org/10.1063/5.0086649, 2022a.

De Florio, M., Schiassi, E., Ganapol, B. D., and Furfaro, R.: Physics-informed neural networks for rarefied-gas dynamics: Poiseuille flow in the BGK approximation, Zeitschrift für angewandte Mathematik und Physik, 73, 126, https://doi.org/10.1007/s00033-022-01767-z, 2022b.

De Florio, M., Kevrekidis, I. G., and Karniadakis, G. E.: AI-Lorenz: A physics-data-driven framework for black-box and gray-box identification of chaotic systems with symbolic regression, Chaos, Solitons & Fractals, 188, 115 538, https://doi.org/10.1016/j.chaos.2024.115538, 2024.

De Florio, M., Zou, Z., Schiavazzi, D. E., and Karniadakis, G. E.: Quantification of total uncertainty in the physics-informed reconstruction of CVSim-6 physiology, Philosophical Transactions A, 383, 20240 221, https://doi.org/10.1098/rsta.2024.0221, 2025.

Demas, N. G., Lorenzo-Martin, C., Luna, R., Erck, R. A., and Greco, A. C.: The Effect of Current and Lambda on White-etch-crack Failures, Tribology International, 189, 108 951, https://doi.org/10.1016/j.triboint.2023.108951, 2023.

Desai, A., Guo, Y., Sheng, S., Phillips, C., and Williams, L.: Prognosis of wind turbine gearbox bearing failures using SCADA and modeled data, in: Annual conference of the PHM society, vol. 12, pp. 1–10, https://doi.org/10.36001/phmconf.2020.v12i1.1292, 2020.

Dong, S. and Li, Z.: Local extreme learning machines and domain decomposition for solving linear and nonlinear partial differential equations, Computer Methods in Applied Mechanics and Engineering, 387, 114 129, https://doi.org/10.1016/j.cma.2021.114129, 2021.

Dwivedi, V. and Srinivasan, B.: Physics informed extreme learning machine (PIELM)–a rapid method for the numerical solution of partial differential equations, Neurocomputing, 391, 96–118, https://doi.org/10.1016/j.neucom.2019.12.099, 2020.

Eftekhari Milani, A., Zappalá, D., Sheng, S., and Watson, S.: Probabilistic prediction of wind turbine remaining useful life using Conformalised Quantile Regression, Manuscript in preparation, unpublished — work in progress. Corresponding author: a.eftekharimilani@tudelft.nl, 2025.

Fabiani, G.: Random projection neural networks of best approximation: Convergence theory and practical applications, SIAM Journal on Mathematics of Data Science, 7, 385–409, https://doi.org/10.1137/24M1639890, 2025.

435   Fabiani, G., Calabrò, F., Russo, L., and Siettos, C.: Numerical solution and bifurcation analysis of nonlinear partial differential equations with extreme learning machines, Journal of Scientific Computing, 89, 44, https://doi.org/10.1007/s10915-021-01650-5, 2021.

Fabiani, G., Bollt, E., Siettos, C., and Yannacopoulos, A. N.: Stability Analysis of Physics-Informed Neural Networks for Stiff Linear Differential Equations, arXiv preprint arXiv:2408.15393, 2024.

Frost, N. E., Marsh, K. J., and Pook, L. P.: Metal fatigue, Courier Corporation, 1999.

440   Galaris, E., Calabrò, F., di Serafino, D., and Siettos, C.: Numerical solution of stiff ordinary differential equations with random projection neural networks, arXiv preprint arXiv:2108.01584, 2021.

Ganaie, M. A., Hu, M., Malik, A. K., Tanveer, M., and Suganthan, P. N.: Ensemble deep learning: A review, Engineering Applications of Artificial Intelligence, 115, 105 151, https://doi.org/10.1016/j.engappai.2022.105151, 2022.

Gawlikowski, J., Tassi, C. R. N., Ali, M., Lee, J., Humt, M., Feng, J., Kruspe, A., Triebel, R., Jung, P., Roscher, R., et al.: A survey of
445   uncertainty in deep neural networks, Artificial Intelligence Review, 56, 1513–1589, https://doi.org/10.1007/s10462-023-10562-9, 2023.

Greco, A., Sheng, S., Keller, J., and Erdemir, A.: Material Wear and Fatigue in Wind Turbine Systems, Wear, 302, 1583–1591, https://doi.org/10.1016/j.wear.2013.01.060, 2013.

Greco, A., Demas, N., Erck, R., Gould, B., Keller, J., Sheng, S., and Guo, Y.: Wind Turbine Drivetrain Reliability, NREL/PR-5000-84029, Golden, CO: National Renewable Energy Laboratory, 2022.

450   Guo, Y., Sheng, S., Phillips, C., Keller, J., Veers, P., and Williams, L.: A Methodology for Reliability Assessment and Prognosis of Bearing Axial Cracking in Wind Turbine Gearboxes, Renewable and Sustainable Energy Reviews, 127, 109 888, https://doi.org/10.1016/j.rser.2020.109888, 2020.

Han, S., Stelz, L., Stoecker, H., Wang, L., and Zhou, K.: Approaching epidemiological dynamics of COVID-19 with physics-informed neural networks, Journal of the Franklin Institute, 361, 106 671, https://doi.org/10.1016/j.jfranklin.2024.106671, 2024.

455   Haus, L., Sheng, S., and Pulikollu, R.: Wind Turbine Major Systems Reliability Trends and Mitigation Strategies, in: Drivetrain Reliability Collaborative Meeting, 2024.

Haus, L., Sheng, S., and Pulikollu, R.: Newer/Larger Wind Turbine Reliability Analysis & Applications, in: Drivetrain Reliability Collaborative Meeting, 2025.

Huang, G.-B., Zhu, Q.-Y., and Siew, C.-K.: Extreme learning machine: theory and applications, Neurocomputing, 70, 489–501,
460   https://doi.org/10.1016/j.neucom.2005.12.126, 2006.

Hüllermeier, E. and Waegeman, W.: Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods, Machine learning, 110, 457–506, https://doi.org/10.1007/s10994-021-05946-3, 2021.

Jensen, O. L., Heuser, L., and Petersen, K. E.: Prevention of 'White Etching Cracks' in Rolling Bearings in Vestas Wind Turbines, in: Conference for Wind Power Drives, 2021.

465   Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., and Yang, L.: Physics-informed machine learning, Nature Reviews Physics, 3, 422–440, https://doi.org/10.1038/s42254-021-00314-5, 2021.

Keller, J., Menck, O., Stammler, M., Pulikollu, R., and Haus, L.: Pitch Bearing Rating and Service Life Analysis for a 1.5-Megawatt Wind Turbine, NREL/TP-5000-94249, Golden, CO: National Renewable Energy Laboratory, 2025.

Kharazmi, E., Cai, M., Zheng, X., Zhang, Z., Lin, G., and Karniadakis, G. E.: Identifiability and predictability of integer- and
470   fractional-order epidemiological models using physics-informed neural networks, Nature Computational Science, 1, 744–753, https://doi.org/10.1038/s43588-021-00158-0, 2021.

Lakshminarayanan, B., Pritzel, A., and Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles, Advances in Neural Information Processing Systems, 30, 2017.

Millevoi, C., Pasetto, D., and Ferronato, M.: A Physics-Informed Neural Network approach for compartmental epidemiological models, PLOS Computational Biology, 20, e1012 387, https://doi.org/10.1371/journal.pcbi.1012387, 2024.

Mortari, D.: Least-squares solution of linear differential equations, Mathematics, 5, 48, https://doi.org/10.3390/math5040048, 2017a.

Mortari, D.: The theory of connections: Connecting points, Mathematics, 5, 57, https://doi.org/10.3390/math5040057, 2017b.

Mortari, D., Johnston, H., and Smith, L.: High accuracy least-squares solutions of nonlinear differential equations, Journal of computational and applied mathematics, 352, 293–307, https://doi.org/10.1016/j.cam.2018.12.007, 2019.

Osorio, J. D., De Florio, M., Hovsapian, R., Chryssostomidis, C., and Karniadakis, G. E.: Physics-Informed machine learning for solar-thermal power systems, Energy Conversion and Management, 327, 119 542, https://doi.org/10.1016/j.enconman.2025.119542, 2025.

Pickering, E., Guth, S., Karniadakis, G. E., and Sapsis, T. P.: Discovering and forecasting extreme events via active learning in neural operators, Nature Computational Science, 2, 823–833, https://doi.org/10.1038/s43588-022-00376-0, 2022.

Psaros, A. F., Meng, X., Zou, Z., Guo, L., and Karniadakis, G. E.: Uncertainty quantification in scientific machine learning: Methods, metrics, and comparisons, Journal of Computational Physics, 477, 111 902, https://doi.org/10.1016/j.jcp.2022.111902, 2023.

Pugno, N., Ciavarella, M., Cornetti, P., and Carpinteri, A.: A generalized Paris' law for fatigue crack growth, Journal of the Mechanics and Physics of Solids, 54, 1333–1349, https://doi.org/10.1016/j.jmps.2006.01.007, 2006.

Rahaman, R. et al.: Uncertainty quantification and deep ensembles, Advances in neural information processing systems, 34, 20 063–20 075, 2021.

Raissi, M., Perdikaris, P., and Karniadakis, G. E.: Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, Journal of Computational physics, 378, 686–707, https://doi.org/10.1016/j.jcp.2018.10.045, 2019.

Raissi, M., Yazdani, A., and Karniadakis, G. E.: Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations, Science, 367, 1026–1030, https://doi.org/10.1126/science.aaw4741, 2020.

Sahli Costabal, F., Yang, Y., Perdikaris, P., Hurtado, D. E., and Kuhl, E.: Physics-informed neural networks for cardiac activation mapping, Frontiers in Physics, 8, 42, https://doi.org/10.3389/fphy.2020.00042, 2020.

Schiassi, E., De Florio, M., D'Ambrosio, A., Mortari, D., and Furfaro, R.: Physics-informed neural networks and functional interpolation for data-driven parameters discovery of epidemiological compartmental models, Mathematics, 9, 2069, https://doi.org/10.3390/math9172069, 2021a.

Schiassi, E., Furfaro, R., Leake, C., De Florio, M., Johnston, H., and Mortari, D.: Extreme theory of functional connections: A fast physics-informed neural network method for solving ordinary and partial differential equations, Neurocomputing, 457, 334–356, https://doi.org/10.1016/j.neucom.2021.06.015, 2021b.

Schiassi, E., De Florio, M., Ganapol, B. D., Picca, P., and Furfaro, R.: Physics-informed neural networks for the point kinetics equations for nuclear reactor dynamics, Annals of Nuclear Energy, 167, 108 833, https://doi.org/10.1016/j.anucene.2021.108833, 2022.

Schleich, F.: Investigation of Cracked Blade Bearing Rings, in: Drivetrain Reliability Collaborative Meeting, 2025.

Shapiro, J.: Pitch Bearing Investigations, in: Drivetrain Reliability Collaborative Meeting, 2017.

Stehly, T. and Duffy, P.: 2020 Cost of Wind Energy Review, NREL/TP-5000-81209, Golden, CO: National Renewable Energy Laboratory, 2021.

Wang, Y. and Dong, S.: An extreme learning machine-based method for computational PDEs in higher dimensions, Computer Methods in Applied Mechanics and Engineering, 418, 116 578, https://doi.org/10.1016/j.cma.2023.116578, 2024.

Wenzel, F., Snoek, J., Tran, D., and Jenatton, R.: Hyperparameter ensembles for robustness and uncertainty quantification, Advances in Neural Information Processing Systems, 33, 6514–6527, 2020.

Wiser, R., Bolinger, M., and Lantz, E.: Assessing Wind Power Operating Costs in the United States: Results from a Survey of Wind Industry Experts, Renewable Energy Focus, 30, 46–57, https://doi.org/10.1016/j.ref.2019.05.003, 2019.

Wiser, R., Millstein, D., Hoen, B., Bolinger, M., Gorman, W., Rand, J., Barbose, G., Cheyette, A., Darghouth, N., Jeong, S., et al.: Land-Based Wind Market Report: 2024 Edition, Tech. rep., Berkeley, CA: Lawrence Berkeley National Laboratory, https://doi.org/10.2172/2434282, 2024.

Yin, M., Zheng, X., Humphrey, J. D., and Karniadakis, G. E.: Non-invasive inference of thrombus material properties with physics-informed neural networks, Computer Methods in Applied Mechanics and Engineering, 375, 113 603, https://doi.org/10.1016/j.cma.2020.113603, 2021.

Zhang, Z., Zou, Z., Kuhl, E., and Karniadakis, G. E.: Discovering a reaction–diffusion model for Alzheimer's disease by combining PINNs with symbolic regression, Computer Methods in Applied Mechanics and Engineering, 419, 116 647, https://doi.org/10.1016/j.cma.2023.116647, 2024.

Zou, Z. and Karniadakis, G. E.: Multi-head physics-informed neural networks for learning functional priors and uncertainty quantification, Journal of Computational Physics, 531, 113 947, https://doi.org/10.1016/j.jcp.2025.113947, 2025.

Zou, Z., Meng, X., and Karniadakis, G. E.: Correcting model misspecification in physics-informed neural networks (PINNs), Journal of Computational Physics, 505, 112 918, https://doi.org/10.1016/j.jcp.2024.112918, 2024a.

Zou, Z., Meng, X., Psaros, A. F., and Karniadakis, G. E.: NeuralUQ: A comprehensive library for uncertainty quantification in neural differential equations and operators, SIAM Review, 66, 161–190, https://doi.org/10.1137/22M1518189, 2024b.