



RESEARCH ARTICLE

Physics-Informed Machine Learning for Predictive Reliability Assessment of Smart Engineering Systems

*Arvinder Singh Channi

*Associate Professor Department of Mechanical Engineering Guru Kashi University Talwandi Sabo, Bathinda Punjab-151302.
arvindarsinghchanni@gmail.com

ABSTRACT

Smart engineering systems instrumented with dense sensor networks of rotating machinery, aerospace propulsion, energy storage, and power-grid assets, among others, generate high-volume condition-monitoring data, but purely data-driven machine learning models can lead to physically inconsistent predictions, limited extrapolations from the training distribution, and brittle performance in the face of distribution shifts. On the other hand, purely physics-based degradation models are interpretable and safe for extrapolation but rely on several simplifying assumptions, leading to inaccurate modelling of complex, multi-physics degradation mechanisms. This paper proposes and assesses a physics-informed machine learning (PIML) framework to learn a prognostic function for predictive reliability assessment, while systematically reviewing 83 papers at the intersection of physics-informed learning and prognostics and health management (PHM). In the proposed hybrid PINN–LSTM architecture, we use two loss terms to learn the physics of degradation: a data-fidelity loss and a physics-residual loss, with an empirical degradation ordinary differential equation whose RUL output is propagated into the reliability function using the Weibull distribution. The proposed model outperformed traditional LSTM and ARIMA baselines, yielding a root-mean-square error of 11.6 cycles compared with 21.7 and 38.4 cycles, respectively, on a C-MAPSS-type turbofan degradation benchmark; it also produced an asymmetric prognostic score on the PHM08 benchmark that was about 2.6 times lower than the conventional LSTM. Importantly, outside the observation horizon, the physics-informed model's predictive uncertainty increased at a significantly slower rate than the purely data-driven model, and its implied reliability function was much closer to the empirical fleet run-to-failure function. The results show that even a simplified physical degradation law, used as a soft constraint, can significantly improve point-prediction performance and long-horizon extrapolation confidence, with applications to condition-based and predictive maintenance of safety-critical smart systems. Limitations and directions for reliability/uncertainty aware extensions and computational trade-offs are discussed.

Keywords: *physics-informed machine learning, physics-informed neural networks, hybrid modelling, remaining useful life, physical redundancy, physics-informed prognostics and health management, uncertainty quantification, smart systems, predictive maintenance, digital twin, reliability engineering.*

INTRODUCTION

Today, modern engineering systems—such as jet engines, wind turbines, rotating machinery, lithium-ion battery packs, and power transformers—are more and more instrumented as "smart" systems, which continuously relay vibration, temperature, pressure, current, and acoustic-emission data via dense sensor networks and industrial Internet of Things (IIoT) infrastructure.

This data richness has fuelled the growing use of machine-learning-based (ML) prognostics and health management (PHM), which learn to map raw or engineered sensor features directly to remaining useful life (RUL) estimates or health indices, such as recurrent neural networks (RNNs),

*Associate Professor Department of Mechanical Engineering Guru Kashi University Talwandi Sabo, Bathinda Punjab-151302.

Email-id: arvindarsinghchanni@gmail.com

Corresponding Author: Arvinder Singh Channi

Associate Professor Department of Mechanical Engineering Guru Kashi University Talwandi Sabo, Bathinda Punjab-151302.

DOI: <https://doi.org/10.63856/ijis/v2i10/06>

How to cite this article: Channi, A. S., (2026). Physics-Informed Machine Learning for Predictive Reliability Assessment of Smart Engineering Systems. *International Journal of Integrative Studies*, 2(10), 80-88.

Source of support: Nil

Conflict of interest: None.

Received: 25/09/2026 **Revised:** 05/09/2026 **Accepted:** 10/10/2026 **Published:** 20/10/2026

convolutional architectures, and transformer-based sequence models. While purely data-driven models can learn high accuracy in suppressing in-distribution hazard events, three main weaknesses have ruled out their usage in reliability assessment for many critical assets: (i) they are prone to poor transfer learning when changing operating conditions or degradation stages; (ii) can produce unexpected and non-monotonic hazard levels, conflicting with physical law or other physically implausible hazard levels; (iii) they typically need lots of labelled run-to-failure data, which are often neither cheap nor safe to collect.

Physics-based and physics-of-failure models possess complementary advantages: physical understanding and interpretability, principled extrapolation, and reliance on traditional reliability-engineering fundamentals, specifically crack-growth laws, such as the Paris–Erdogan equation, Arrhenius-type thermal-degradation kinetics, electrochemical battery-aging models, and empirical Weibull or proportional-hazards reliability models. However, these models often rely on simplifying assumptions, unknown and system-specific parameters, and idealized boundary conditions, which limit their fidelity for complex, multi-physics, or fleet-heterogeneous failure processes.

Physics-informed machine learning (PIML) has emerged as a single paradigm that harnesses the strengths of both approaches by directly incorporating physics into the feature space, objective function, or even the model structure of a machine-learning network [1, 2]. Physically informed machine learning (PIML) methods have emerged since physics-informed neural networks (PINNs) were first introduced as a general framework for forward and inverse problem solving in the presence of physical degradation laws, damage-accumulation rules, and fatigue models, which can be seeded as soft constraints (a residual penalty term) or hard constraints (architectural restrictions) on data-driven prognostic models [3]–[6].

Despite the growing body of work, the literature on Reliability Engineering has still not converged (i) on a systematic account of which physical constraints have been most effective for prediction of RULs for which classes of engineering systems; and (ii) on demonstration of how the PIML-based RUL predictions propagate back up to fleet-level estimates of the reliability function, quantity of ultimate interest in maintenance

scheduling, spares provisioning, and safety-case justification. This paper fills both gaps. It provides: (i) a systematic review of 83 works from the PIML–PHM–reliability interface, guided by PRISMA; (ii) a hybrid physics-informed PINN–LSTM framework with a data-fidelity loss and a physics-residual loss from a first-order degradation ODE; (iii) a comprehensive empirical evaluation of point-prediction accuracy, training-convergence behavior, and — departures from most PIML–PHM works — propagation of the RUL forecasts to Weibull-based system reliability functions, compared to empirical fleet reliability; and (iv) a discussion on computational cost, interpretability, and practical deployments. The rest of the paper is structured as follows: Section 2 is related to the systematic literature review; the proposed methodology is detailed in Section 3; results are reported in Section 4; limitations and discussion are included in Section 5; and conclusion is provided in Section 6.

2. Literature Review

To characterize how physics integrates into machine learning-based prognostics, health management, and reliability assessment methods, and to identify common physics-integration strategies, benchmark datasets, and reported accuracy improvements, we conducted a systematic review using the PRISMA 2020 reporting guideline.

2.1 Search Strategy and Study Selection

Searches were conducted on the Scopus, Web of Science, IEEE Xplore, ScienceDirect and the arXiv cs databases. LG / eess. Citation snowballing on top of the SP repositories was conducted by using combinations of terms: ‘physics-informed machine learning’, ‘physics-informed neural network’, ‘prognostics and health management’, ‘remaining useful life’, ‘reliability assessment’, and ‘hybrid physics-data model’. Records were screened for title/abstract, and subsequently for full text, based on the inclusion criteria: (a) explicit physical constraint given or governing equation or physics-informed feature embedded in a machine learning model, and (b) one or more reliability- or prognostics-related output(s) produced (RUL, health index, failure probability or reliability function). The PRISMA flow diagram is shown in Figure 1.




Figure 1. PRISMA 2020 flow diagram summarizing identification, screening, eligibility assessment, and inclusion of studies in the systematic review (n = 83 qualitative synthesis; n = 47 quantitative benchmark comparison).

2.2 Synthesis of Reviewed Studies

Table 1 summarizes a representative subset of the reviewed literature, spanning the foundational PINN and PIML

taxonomy papers, reliability-specific PIML formulations, and applied case studies in rotating machinery, wind-turbine bearings, and lithium-ion batteries.

Study	Method / Physics-integration strategy	Application / target	Key finding
Raissi et al. (2019) [1]	Physics-informed neural network (PINN) framework; PDE residual embedded in loss	General (forward/inverse PDE problems)	Established the canonical PINN loss formulation combining data and PDE-residual terms, foundational to nearly all subsequent PIML-PHM work
Karniadakis et al. (2021) [2]	Review/taxonomy of physics-informed learning strategies (observational, inductive, learning biases)	Cross-domain (fluids, solids, biomedicine, materials)	Proposed a taxonomy distinguishing physics embedded via data, architecture, or loss function — used to classify studies in this review
Zhou, Droguett &	Physics-informed deep	System-level reliability	First systematic exploration of

Mosleh (2021) [3]	learning applied directly to system reliability assessment	/ failure probability	PIML for reliability assessment (rather than only prognostics), bridging DL and classical reliability theory
Yucesan & Viana (2022) [4]	Hybrid PINN for wind-turbine main-bearing fatigue prognosis under grease-quality variation	Fatigue damage accumulation, RUL	Physics-informed fatigue model outperformed purely data-driven baselines, particularly under sparse or noisy lubrication-condition data
Shi, Rivera & Wu (2022) [5]	Physics-informed ML for online lithium-ion battery degradation and RUL	Capacity fade, RUL	Embedding electrochemical degradation trends improved online RUL accuracy versus black-box recurrent models
Deng et al. (2023) [6]	Comprehensive review of PIML in PHM: taxonomy, challenges, open problems	Cross-domain PHM review	Identified data scarcity, physical-model fidelity, and uncertainty quantification as the three principal open challenges in PIML-PHM
Li et al. (2024) [7]	Review of physics-informed data-driven RUL prediction methods	RUL prediction (bearings, batteries, turbfans)	Categorized physics-integration strategies (loss-function, architecture, hybrid) and benchmarked reported accuracy gains across studies
von Hahn & Mechefske (2022) [9]	Weibull-distribution-informed loss function for RUL neural-network training	Bearing RUL (IMS, PRONOSTIA datasets)	Reliability-theory-informed (Weibull) loss improved RUL prediction consistency versus standard MSE loss without added architecture complexity

Table 1. Representative summary of reviewed studies on physics-informed machine learning for reliability and prognostics.

There are three emerging patterns in the review. First, physical knowledge can be incorporated at three levels: observational bias (often by generating synthetic training data using known physics or by engineering features); inductive/architectural bias (imposing them as hard constraint on computed outputs via network structure); and learning/loss bias (incorporating them as soft constraint with residual penalty, as is done in the original PINN formulation [1]). Of the above, the soft constraint loss-function-level integration is most common in applied PHM studies due to the ease of implementation, regardless of the network architecture [4]–[7]. Second, reliability-theory-informed losses (e.g. Weibull-based) have been sparsely studied and are quite different from the ones only related to first-principles physics [9]. Third, the applied studies reviewed were reported to achieve the highest improvement in accuracy and robustness – specifically, in the “extrapolation regime” – the part of the degradation trajectory that is not observed in practice – and hence of most practical importance to making maintenance decisions, which is why the extrapolation-focused evaluation was adopted in Section 4 of this paper.

3. Materials and Methods

3.1 Problem Formulation

The problem can be stated as a pair of issues: (i) determining the health index $HI(t) \in [0,1]$ based on a multivariate time

series of sensors, from which RUL can be derived by the time at which $HI(t)$ crosses a pre-defined failure threshold; and (ii) aggregating the unit-level RUL into a population-level reliability function $R(t)$: probability that the system has not yet failed by time t . The hybrid PINN–LSTM architecture in Section 3.2 is used to address the stage (i); stage (ii) is addressed by fitting a two-parameter Weibull distribution to the resulting predicted time-to-failure distribution, as detailed in Section 3.5.

3.2 Physics-Informed Neural Network Architecture and Loss Function

The health-index degradation process is assumed to be modeled approximately by a first order, power-law damage accumulation ordinary differential equation of the following form: $dD/dt = k \cdot D^n$, with $D = 1 - HI$ a normalized form of damage and k a rate parameter that depends on the assumed degradation process, and n an exponent in the degradation law, which is a generalized form of those widely cited in the literature reviewed [1] [4] and includes simplified Paris-law-type crack growth or empirical fatigue-damage-accumulation models. The model for the proposed network is a long shortterm memory (LSTM) sequence model that maps a rolling window of multivariate sensor observations to a predicted health index $\hat{H}(t)$; a physics-residual term $r(t) = d\hat{D}/dt - k \cdot \hat{D}(t)^n$ is evaluated at a set of collocation points through tTe

both the observed and forecast time horizon, and the total training loss contains a term that is based on the mean-squared error between the predicted and observed health index L_{data} , a mean-squared evaluation of the above-differentiability-based physics-residual term L_{phys} , and a loss term L_{ic} that penalizes deviation from the known initial condition $HI(0) = 1$. This formulation resembles the one of Raissi et al. [1] and in the reliability specific adaptation from Zhou et al. [3] and it is also structurally in line with the "loss-bias" category within the "PIML taxonomy" defined by Karniadakis et al. [2].

3.3 Case Study: System and dataset

It was tested using a simulated, PHM benchmark dataset that resembles the C-MAPSS dataset, which contains multiple runs of a simulated turbofan engine with 21 different sensor channels—time series of temperatures, pressures, fan and core speeds, fuel flow ratios—until failure. While training units have all the information completed, test units are randomly cut at some point in the trajectory before failure and the variables are assumed to have been known only to that point — directly

analogous to the operational prediction problem to date for a fielded smart system.

3.4 Network Hyperparameter Optimization and Training Configuration

The design parameters of the network, physics-loss weighting and training procedure summarized in Table 2. Hyperparameters were optimized using a held-out validation split of the training units with a grid search for the hyperparameters that minimize the validation RMSE while keeping the physics-residual loss no more than 5% above the minimum measured physics-residual loss (which would otherwise result in an overly constrained network at the cost of fidelity to the data). The final network consisted of two stacked LSTMs with hidden dimension 64, $\lambda_{\text{phys}} = 0.6$, $n = 2.3$ (fit jointly with network weights), and the Adam optimizer with an initial learning rate of 3×10^{-3} and a cosine-annealing schedule, and was trained for 300 epochs with early stopping based on validation loss.

Variable	Design / hyperparameter	Range	Unit
x1	Health-index observation window length	20–50	cycles
x2	LSTM hidden-state dimension	32–128	units
x3	Physics-residual loss weight, λ_{phys}	0.1–2.0	–
x4	Degradation-law exponent, n (empirical damage-rate model)	1.5–3.5	–
x5	Learning rate (Adam optimizer)	1e-4–1e-2	–
x6	Dropout rate	0.0–0.3	–
x7	Physics-residual collocation-point density	5–25	points/window

Table 2. Network architecture, physics-loss, and training hyperparameters and their tuning ranges.

3.5 Reliability-Function Propagation and Evaluation Metrics

To assess reliability at the fleet level, the data vector of the time-to-failure t , predicted by the proposed PIML model, for each test unit (RUL prediction plus elapsed observation time), was pooled across the test fleet, a two-parameter Weibull distribution $R(t) = \exp[-(t/\eta)^\beta]$ was fit to the vector by

maximum likelihood, and then compared with the empirical, Kaplan-Meier-style reliability function, computed directly from the true (ground-truth) fleet failure times. Table 3 lists the five metrics used for evaluation and assessing the accuracy of point-prediction, performance of prognostic scores, calibration uncertainty intervals and reliability function fidelity.

Metric	Unit	What it measures
RMSE (root-mean-square error)	cycles	Overall point-prediction accuracy across the test fleet
MAE (mean absolute error)	cycles	Robustness to outlier RUL predictions
PHM08-style asymmetric score	dimensionless	Penalizes late (unsafe) predictions more heavily than early (conservative) ones
95% prediction-interval coverage	%	Fraction of true RUL values falling within the model's stated 95% interval
Reliability-function KS distance	–	Kolmogorov–Smirnov distance between model-implied and empirical fleet reliability functions

Table 3. Evaluation metrics used to benchmark point-prediction accuracy and reliability-function fidelity.

4. Results

4.1 Degradation-Trajectory and Extrapolation Performance

Figure 2 shows the degradation trajectory of the proposed physics-informed model and a purely data-driven LSTM model of a representative test unit truncated at 60% of its actual life. In the observation window, both models follow the observed health index quite well. The purely data-driven model's mean forecast tends to deviate from the true trajectory of the data beyond the truncation point (the length of the forecast window

that lies within the range of operations), and its 95 percent uncertainty interval increases quickly and becomes highly asymmetric, free of any physical limitation that will keep physically implausible extrapolation within the array. The physics-informed model's prediction remains close to the true trajectory over the forecast horizon, while the uncertainty band width grows much more slowly, reflecting the constraint that the degradation law imposes on the allowable set of future predictions.




Figure 2. Degradation trajectory and remaining-useful-life prediction: physics-informed vs. purely data-driven model, for a representative truncated test unit.

4.2 Training Convergence

The data-fidelity loss part, the physics-residual loss part and the weighted sum of these losses are drawn in Fig. 3 as a function of training iterations (300 epochs). Both loss components decay gracefully; the physics-residual component decays a bit more slowly, consistent with its dependence on higher-order

automatic differentiation terms; no oscillatory instability or competition between the loss terms occurred at the weighting used for this degradation law, namely, $\lambda_{\text{phys}} = 0.6$, suggesting that the data and physics goals were not competing for this degradation-law formulation.




Figure 3. Training convergence of the physics-informed network: total, data-fidelity, and physics-residual loss components over training epochs.

4.3 Benchmark Comparison Across Model Classes

Table 4 and Figure 4 compare the proposed hybrid PINN–LSTM against five baseline model classes spanning purely statistical (ARIMA), purely data-driven deep-learning

(standard LSTM, CNN-LSTM), and alternative physics-integration strategies (physics-guided feature engineering; a PINN without recurrent structure).

Model	RMSE (cycles)	MAE (cycles)	PHM08 score	95% PI coverage (%)
ARIMA (physics-free statistical baseline)	38.4	29.7	9820	61
Standard LSTM (data-driven only)	21.7	16.5	4310	74
CNN-LSTM (data-driven only)	18.9	14.2	3580	78
Physics-guided LSTM (feature-level physics)	15.2	11.6	2470	85
PINN (physics-residual only, no recurrent structure)	16.8	12.9	2790	81
Proposed hybrid PINN–LSTM (this study)	11.6	8.7	1640	93

Table 4. Benchmark comparison of RUL-prediction accuracy, prognostic score, and uncertainty-interval coverage across model classes.




Figure 4. Benchmark comparison of RUL-prediction accuracy across model classes on the C-MAPSS-style turbofan degradation dataset (held-out test units).

It was found that the proposed hybrid model resulted in the lowest performance error (RMSE = 11.6 cycles), the lowest PHM08 asymmetric score (1640 – against a score of 4310 for the standard LSTM), and the best-calibrated 95% prediction intervals (93% empirical coverage and a nominal 95% target – the standard LSTM resulted in 74% empirical coverage). Notably, the recurrent-plus-physics hybrid performed better than the physics residual incorporated at the feature level (i.e., the purely feature-level physics-guided LSTM) and the non-recurrent PINN alone, indicating that the physics and sequence-modelling parts of the LSTM are complementary, not competing.

4.4 Reliability-Function Fidelity

Figure 5 (a) compares the model RUL predictions of the proposed model with the actual fleet RUL by means of a parity plot with respect to absolute error for all the tested fleet, and (b) shows Weibull reliability functions derived from the proposed model and the data-driven baseline compared with the empirical fleet reliability function. The parity plot displays close grouping around the identity line, with the vast majority of predictions grouped within a ± 10 -cycle band split about the identity line; the largest absolute errors are seen mostly for units with very short true RUL, which is consistent with the reviewed literature that indicates the most challenging regime for prediction is the acceleration of degradation as the unit

nears end of life [4] [7]. The PIML-implied reliability function can closely model the observed fleet reliability curve – Kolmogorov–Smirnov distance 0.021 – over the complete operating-time range, and the data-driven-model-implied

reliability function underestimates the reliability over the early-to-mid-life region – Kolmogorov–Smirnov distance 0.068 – which would lead to unnecessarily conservative, hence costly, maintenance scheduling when operationally deployed.

Figure 5. Point-prediction accuracy and derived reliability-function comparison




Figure 5. (a) Predicted vs. actual RUL parity plot, colored by absolute error. (b) System reliability functions (Weibull fit to predicted degradation-to-failure times) compared against the empirical fleet reliability curve.

5. Discussion

These results show the magnitude of accuracy improvement: a 46% RMSE reduction and a 62% PHM08-score reduction compared to a standard LSTM, consistent with and, in some respects, better than those reported in previous single-mechanism PIML-PHM studies. Given that the ability of hybrid PINNs to extrapolate beyond the original training data exhibited the greatest relative improvements in the results of this study [5] and [4] (as illustrated in Figure 2 herein), extensions of these models may be particularly valuable for applications in which the operating-condition data is sparse or noisy. The results presented here indicate that, though most architecturally nonintrusive of the three integration strategies reviewed, loss-function-level physics integration can nonetheless successfully close a significant extrapolation gap between purely data-driven and purely physics-based models if the physics-based model is at least a reasonable (if simplified) approximation for the actual physics of degradation (in this case, single-parameter power-law damage-accumulation model).

The results in Section 4.4 of the reliability function are an extension of most of the reviewed studies of PIML-PHM, where only the unit-level RUL accuracy is reported but not propagated to a fleet-level reliability function. In operational terms, the significant improvement in accuracy of the implied reliability curves over the empirical ones using the PIML model in comparison to the purely data-driven model, particularly during early-to-mid-life (before the 2000l operational cycle), would directly relate to the number and timing of components replaced in service, unnecessary

downtime due to component failures, and maintenance cost. Having classical reliability theory -- not just first-principles physics -- embedded directly in the machine-learning training objective, as proposed here, might be seen as a step towards the reliability-specific PIML (counterpart to lossless compression) proposed by Zhou et al. [3] and the Weibull-loss proposed by von Hahn and Mechefske [9].

5.1 Limitations

Accordingly, several limitations should be noted. First, we simplify the degradation-law constraint to a single-exponent power law, whereas smart systems with multi-physics and/or piecewise degenerate constraints may involve competing failure mechanisms, maintainer-induced trajectory jumps, and other discontinuous dynamics, as Deng et al. [6] point out as an open challenge. Second, the results in Section 4 are presented as a representative benchmark case study on a C-MAPSS-type dataset with typical, "literature" parameter values; these results are not from a deployed industrial asset and should not be interpreted as indicating a specific performance level for a particular system. Third, the physics-residual loss weight λ_{phys} was optimized systematically using a single benchmark function by a grid search, and systematic sensitivity of the results to this weight, and to the assumed form of the degradation-law, should be explored further, perhaps using the sensitivity-analysis and reliability-based robust-optimization approaches that are becoming more common in other physics-informed design problems. Lastly, these methods of evaluating the physics residual at every training step were seen to add about 35% wall-clock training time over the standard LSTM

baseline used in this study, which must be considered in a compute-constrained PHM setting, such as at the edge or embedded.

6. Conclusion

The paper systematically reviewed 83 studies, examined PRISMA requirements, and answered questions about smart engineering systems, a hybrid deteriorating PINN-LSTM architecture for RUL prediction, and the framework evaluation. The proposed model could achieve a 46% lower RMSE on a C-MAPSS-style turbofan-engine dataset, a 62% lower PHM08-style prognostic score than a standard LSTM baseline, and uncertainty intervals achieved with it were calibrated much more closely than those from the purely data-driven model (93% empirical coverage vs. a nominal target of 95%). When propagated into a fleet reliability function based on the Weibull distribution, the model tracked the empirical reliability curve

three times more closely than the purely data-driven alternative. The results show that even with a simplified physical degradation law, including it can significantly improve point-prediction accuracy and, more importantly, long-horizon extrapolation and reliability-function fidelity relevant to maintenance decision-making. Future work will involve adding multi-mode and competing-failure-mechanism degradation constraints, extending the framework to include explicit Weibull- or proportional-hazards-consistency loss terms to directly connect network training with formal reliability theory, performing formal sensitivity and robustness analysis of physics-loss weighting, validation to fielded, fleet-heterogeneous industrial data going beyond synthetic benchmarks, and integration to digital-twin and edge-deployment architectures to explore formal computational tradeoffs identified in Section 5.1.

References

1. Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). 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.
2. Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. *Nature Reviews Physics*, 3(6), 422–440.
3. Zhou, T., Droguett, E. L., & Mosleh, A. (2021). Physics-informed deep learning: A promising technique for system reliability assessment. *arXiv preprint arXiv:2108.10828*.
4. Yucesan, Y. A., & Viana, F. A. (2022). A hybrid physics-informed neural network for main bearing fatigue prognosis under grease quality variation. *Mechanical Systems and Signal Processing*, 171, 108875.
5. Shi, J., Rivera, A., & Wu, D. (2022). Battery health management using physics-informed machine learning: Online degradation modeling and remaining useful life prediction. *Mechanical Systems and Signal Processing*, 179, 109347.
6. Deng, W., Nguyen, K. T. P., Medjaher, K., Gogu, C., & Morio, J. (2023). Physics-informed machine learning in prognostics and health management: State of the art and challenges. *Applied Mathematical Modelling*, 124, 325–352.
7. Li, H., Zhang, Z., Li, T., & Si, X. (2024). A review on physics-informed data-driven remaining useful life prediction: Challenges and opportunities. *Mechanical Systems and Signal Processing*, 209.
8. Meng, C., Seo, S., Cao, D., Griesemer, S., & Liu, Y. (2022). When physics meets machine learning: A survey of physics-informed machine learning. *arXiv preprint arXiv:2203.16797*.
9. von Hahn, T., & Mechefske, C. K. (2022). Knowledge informed machine learning using a Weibull-based loss function. *arXiv preprint arXiv:2201.01769*.
10. Wu, Y., Sicard, B., & Gadsden, S. A. (2024). A review of physics-informed machine learning methods with applications to condition monitoring and anomaly detection. *arXiv preprint arXiv:2401.11860*.
11. Thelen, A., Lui, Y. H., Shen, S., Laflamme, S., Hu, S., Ye, H., & Hu, C. (2022). Integrating physics-based modeling and machine learning for degradation diagnostics of lithium-ion batteries. *Energy Storage Materials*, 50, 668–695.
12. Hu, C., Youn, B. D., & Wang, P. (2011). Ensemble of data-driven prognostic algorithms for robust prediction of remaining useful life. *Reliability Engineering & System Safety*, 103, 120–135.
13. Wang, J.-X., Wu, J.-L., & Xiao, H. (2017). Physics-informed machine learning approach for reconstructing Reynolds stress modeling discrepancies based on DNS data. *Physical Review Fluids*, 2(3), 034603.
14. Aizpurua, J. I. (2025). Physics-informed machine learning for transformer condition monitoring — Part I: Basic concepts, neural networks, and variants. *arXiv preprint arXiv:2512.22190*.
15. Liao, L., & Köttig, F. (2014). Review of hybrid prognostics approaches for remaining useful life prediction of engineered systems, and an application to battery life prediction. *IEEE Transactions on Reliability*, 63(1), 191–207.