




ELSEVIER

Contents lists available at ScienceDirect

# Reliability Engineering and System Safety

journal homepage: [www.elsevier.com/locate/ress](http://www.elsevier.com/locate/ress)

## Review

### From physics to machine learning and back: Part II - Learning and observational bias in prognostics and health management (PHM)

Olga Fink <sup>1</sup>\*, Ismail Nejjar <sup>1</sup>, Vinay Sharma <sup>1</sup>, Keivan Faghieh Niresi <sup>1</sup>, Han Sun <sup>1</sup>, Hao Dong <sup>1</sup>, Chenghao Xu <sup>1</sup>, Amaury Wei <sup>1</sup>, Arthur Bizzi <sup>1</sup>, Raffael Theiler <sup>1</sup>, Yuan Tian <sup>1</sup>, Leandro Von Krannichfeldt <sup>1</sup>, Zhan Ma <sup>1</sup>, Sergei Garmaev <sup>1</sup>, Zepeng Zhang <sup>1</sup>, Mengjie Zhao <sup>1</sup>, Kevin Steiner <sup>1</sup>, Yusuf Kesmen <sup>1</sup>

*Intelligent Maintenance and Operations Systems Lab, EPFL, Lausanne, Switzerland*

#### ARTICLE INFO

##### Keywords:

Prognostics and health management  
Physics-informed machine learning  
Learning bias  
Observational bias  
PINNs  
Virtual sensing  
Data augmentation  
Reinforcement learning  
Domain generalization

#### ABSTRACT

Prognostics and Health Management (PHM) ensures the reliability, safety, and efficiency of complex engineered systems by enabling fault detection, anticipating equipment failures, and optimizing maintenance activities throughout an asset's lifecycle. However, real-world PHM presents persistent challenges: sensor data is often noisy or incomplete, available labels are limited, and degradation behaviors and system interdependencies can be highly complex and nonlinear. Physics-informed machine learning has emerged as a promising approach to address these limitations by embedding physical knowledge into data-driven models. This review examines how incorporating learning and observational biases through physics-informed modeling and data strategies can guide models toward physically consistent and reliable predictions. Learning biases embed physical constraints into model training through physics-informed loss functions and governing equations, or by incorporating properties like monotonicity. Observational biases influence data selection and synthesis to ensure models capture realistic system behavior through virtual sensing for estimating unmeasured states, physics-based simulation for data augmentation, and multi-sensor fusion strategies. The review then examines how these approaches enable the transition from passive prediction to active decision-making through reinforcement learning, which allows agents to learn maintenance policies that respect physical constraints while optimizing operational objectives. This closes the loop between model-based predictions, simulation, and actual system operation, empowering adaptive decision-making. Finally, the review addresses the critical challenge of scaling PHM solutions from individual assets to fleet-wide deployment. Fast adaptation methods, including meta-learning and few-shot learning, are reviewed alongside domain generalization techniques that ensure robust performance across diverse assets and fleets. These advances are essential for moving PHM from isolated applications to broad, impactful deployment across real-world contexts. For each approach, we discuss its strengths, limitations, and relevance for different PHM tasks.

## 1. Introduction

Prognostics and Health Management (PHM) systems aim to ensure the long-term reliability, availability, and safety of complex industrial and infrastructure assets by predicting system degradation and supporting optimal maintenance decisions [1]. Increasingly, PHM leverages machine learning (ML) techniques for fault detection [2], diagnostics [3], and Remaining Useful Life (RUL) prediction [4]. However, practical deployment of PHM continues to face persistent challenges: industrial data is often noisy or incomplete, labels are limited, and samples of faulty or

degraded conditions may be partly or entirely missing. In addition, underlying degradation processes are frequently complex and nonlinear. These obstacles make it difficult for purely data-driven models to generalize beyond their training conditions or produce physically consistent predictions. Overcoming these challenges requires approaches that combine the strengths of data-driven learning with insights grounded in the physics of the underlying systems.

**Physics-informed machine learning (PIML)** [5] has emerged as a powerful strategy to address these challenges by embedding physical knowledge into ML model structures and training. By integrating

\* Corresponding author.

E-mail address: [olga.fink@epfl.ch](mailto:olga.fink@epfl.ch) (O. Fink).

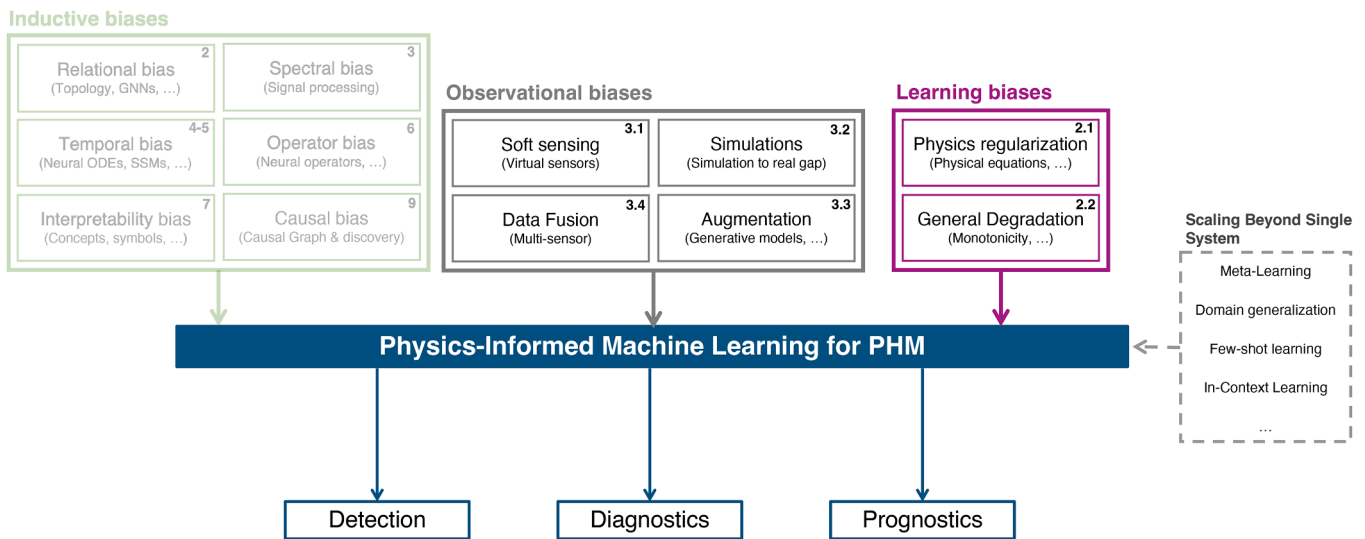
<sup>1</sup> Equal Contributions.

<https://doi.org/10.1016/j.ress.2026.112376>

Received 23 August 2025; Received in revised form 31 January 2026; Accepted 5 February 2026

Available online 13 February 2026

0951-8320/© 2026 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).



**Fig. 1.** Physics-Informed Machine learning framework for PHM. Part I addressed inductive biases that embed structural assumptions into model architectures. Part II (this article) examines observational biases and learning biases that shape data representation and training dynamics. The framework extends to scaling approaches that enable fast adaptation and deployment across diverse systems.

physical laws directly into the structure or training of machine learning algorithms, PIML approaches ensure that model predictions inherently respect known physical constraints, even when observations are limited. Within this context, three primary types of bias have been identified as essential for steering models toward physically consistent and reliable solutions:

1. **Inductive biases** are prior assumptions explicitly embedded into model architectures
2. **Observational biases** influence the model through constraints imposed on the selection, representation and sampling of input data, ensuring the data reflects relevant physical realities.
3. **Learning biases** shape the training process itself, arising from the choice of algorithms, optimization objectives, and regularization strategies.

This paper is the second part of a comprehensive review of physics-informed machine learning for PHM:

- Part I explored inductive biases, architectural design strategies that embed physical knowledge directly into model structures [6].
- Part II (this article) focuses on learning and observational biases, which influence the training process and data representation, as well as on scaling approaches that support robust deployment across diverse operational contexts.

**Fig. 1** illustrates how learning and observational biases complete the physics-informed framework for PHM. This paper presents a comprehensive perspective that begins with (1) the role of learning and observational biases in guiding models to learn physics-consistent system behavior from data, to (2) demonstrates how these physics-informed predictions enable intelligent maintenance decisions through reinforcement learning, and (3) explores strategies for scaling these approaches to manage and optimize entire fleets of assets.

For completeness, the literature selection criteria for *Part II* are summarized in **Table 1**. This table provides an overview of the representative search keywords, reference periods, and primary repositories consulted for each section, ensuring transparency and reproducibility in how the reviewed works were identified. As in Part I, our intention is exemplary rather than exhaustive: the table highlights influential and recent contributions that best illustrate the methodological landscape of observational and learning biases in PHM, guid-

ing readers through the breadth of developments covered in this review.

Across both parts of this review, we adopt a consistent structure: we first move from physics to machine learning by embedding physical laws into data-driven models, and then close the loop by returning from machine learning to physics through the derivation of symbolic physical equations in Part I, and through interaction with the physical world via feedback and control in Part II.

**Learning biases** (see **Section 2**) are introduced through the careful selection of loss functions, constraints, and optimization algorithms, guiding machine learning models toward solutions that align with underlying physical laws. Unlike the architectural constraints of inductive biases examined in Part I, learning biases are explicitly imposed during model training to ensure that learned representations and functions remain consistent with known system behaviors. In PHM applications, learning biases play a critical role in overcoming the fundamental challenge of data scarcity. Industrial datasets are often limited in size, coverage, or diversity, making it difficult for machine learning models to learn robust, generalizable patterns. By embedding physical laws and constraints directly into the learning process, learning biases enable models to compensate for missing or sparse data, effectively guiding them toward solutions that remain faithful to the underlying physics even when observations are incomplete. Without such guidance, purely data-driven optimization can lead models to overfit limited training data, resulting in predictions that violate basic physical principles [7] or producing unrealistic degradation trajectories.

In this review, we examine two primary mechanisms for introducing learning biases in PHM. Physics-Informed Neural Networks (PINNs) (see **Section 2.1**), incorporate governing differential equations directly into the loss function, ensuring that model solutions satisfy conservation laws and system dynamics. As a complementary approach, General Degradation Dynamics Informed Prognostics Methods (see **Section 2.2**) offer a more flexible alternative to physics-based or domain-specific approaches by embedding general physical properties, such as monotonicity and irreversibility, that characterize degradation processes across diverse systems.

**Observational biases** (see **Section 3**) capture the patterns and distributions inherent in the training data. These biases can be leveraged not only by curating datasets that both adhere to physical principles and comprehensively represent the system's operational states, but also by enriching the data through virtual observations and simulated

**Table 1**  
Summary of literature selection criteria across all sections (Part II).

Section	Representative search keywords	Period of references covered	Primary repositories / sources consulted
<b>2.1 Physics-Informed Neural Networks (Learning Bias)</b>	physics-informed neural networks; PINN PHM; PDE-constrained learning; physics-regularized loss; physics-constrained optimization; PINN industrial prognostics; stiff PDE PINN; degradation PDE learning	2017–2025 (from original PINNs to recent PHM-specific and hybrid PINN developments)	NeurIPS, ICLR, ICML; IEEE Xplore (T-Reliab., T-ASE, T-NNLS); Elsevier (RESS, MSSP); SpringerLink
<b>2.2 General Degradation Dynamics (Learning Bias)</b>	monotonic degradation modeling; trendability constraints; general degradation dynamics; irreversible degradation learning; physics-inspired RUL constraints; health indicator monotonicity learning; degradation-informed prognostics	2010–2025 (from classical HI models to modern differentiable monotonicity/convexity constraints in PHM)	IEEE Xplore (T-Reliab., T-IM, T-IE); Elsevier (RESS, MSSP); PHM Society papers; SpringerLink
<b>3.1 Soft Sensing (Virtual Sensing)</b>	soft sensor PHM; virtual sensing state estimation; Kalman-based soft sensing; model-driven soft sensors; hybrid soft sensor; observer-based PHM; virtual sensor RUL; industrial soft sensing review	2000–2025 (from classical Kalman/observer formulations to deep-learning soft sensors and hybrid virtual sensing)	IEEE Xplore (T-IM, T-IE, T-Reliab.); Elsevier (RESS, Chem Eng. Res., MSSP); SpringerLink; MDPI Sensors
<b>3.2 Physics-Based Simulation / Sim-to-Real</b>	simulation-to-real PHM; synthetic data PHM; physics-based simulation prognostics; FEM/CFD PHM simulation; sim-to-real domain gap; synthetic degradation trajectory generation; transfer learning PHM; domain adaptation PHM	2005–2025 (from early synthetic-data PHM studies to modern domain-invariant sim-to-real strategies)	IEEE Xplore; Elsevier (MSSP, RESS); ACM DL; NeurIPS / ICLR domain adaptation papers; PHM Society datasets
<b>3.3 Generative Models for Augmentation</b>	PHM generative models; GAN PHM; VAE PHM; physics-informed generative modeling; rare-failure generation; synthetic fault signal generation; degradation manifold modeling; physics-guided GAN	2014–2025 (from early GAN/VAE applications to physics-informed, data-efficient generative models)	NeurIPS, ICML, ICLR; IEEE Xplore (T-II, T-IM); Elsevier (MSSP, RESS); MDPI Sensors
<b>3.4 Data Fusion (Multimodal Observational Bias)</b>	multimodal PHM fusion; heterogeneous sensor fusion; vibration-acoustic-thermal fusion; multimodal forecasting; multi-sensor fusion PHM; cross-modal degradation modeling; PIML multimodal fusion	2010–2025 (from classical feature-level fusion to recent deep multimodal fusion and PIML-aware fusion architectures)	IEEE Xplore (T-Reliab., T-IM, T-II); Elsevier (MSSP, RESS); SpringerLink; ACM DL; MDPI Sensors
<b>4. Reinforcement Learning for PHM Decision-Making</b>	reinforcement learning PHM; health-aware RL; physics-informed RL; maintenance policy learning; RL degradation control; safe RL industrial systems; model-based RL PHM	2016–2025 (from standard deep RL to physics-aware, safe, and maintenance-oriented RL for PHM)	NeurIPS, ICML, ICLR; IEEE Xplore (T-AI, T-Reliab.); Elsevier (RESS); SpringerLink
<b>5.1 Meta-Learning and Few-Shot Adaptation</b>	meta-learning PHM; few-shot PHM; fast adaptation prognostics; MAML prognostics; meta-RUL; fleet adaptation PHM; cross-asset transfer PHM	2017–2025 (from early meta-learning to cross-asset PHM adaptation and fleet-level fast learning)	NeurIPS, ICML, ICLR; IEEE Xplore; Elsevier (RESS, MSSP); PHM Conference papers
<b>5.2 Domain Generalization for PHM</b>	domain generalization PHM; DG prognostics; out-of-distribution PHM; invariant learning PHM; cross-domain fault detection; distribution shift PHM	2013–2026 (from early DG/OOD theory to PHM-specific DG under physical constraints)	NeurIPS, ICLR, ICML; IEEE Xplore; Elsevier (RESS, MSSP); SpringerLink; DG/OOD benchmarks

samples from conditions that were not originally observed. Conceptually, observational biases offer the most straightforward way to embed physical knowledge into models, exposing them to physically meaningful data, thereby improving robustness. However, PHM applications present unique challenges for data collection: sensors may be sparse or expensive, or impossible to install in certain locations; some failure modes are rare; and specific system states may be impossible to observe directly. These constraints often result in datasets with significant gaps and biases, which can mislead purely data-driven models. Physics-informed observational biases address these limitations through four key strategies.

Virtual sensing (see [Section 3.1](#)) leverages known physical relationships to estimate unmeasured variables or reconstruct missing data using physics-based models, effectively expanding the observable state space of the system. Physics-based simulations (see [Section 3.2](#)) generate synthetic training data using first-principles models or numerical solvers, enabling the creation of representative samples for rare events and a

wider range of operating conditions. Physics-informed data augmentation (see [Section 3.3](#)) increases the dataset diversity by generating new data while strictly preserving physical constraints. Finally, multi-sensor data fusion (see [Section 4](#)) integrates complementary sensor modalities to construct richer and more complete representations of system behavior. Through these combined methods, models are exposed to data that more closely mirrors the true physical mechanisms of degradation. [Table 2](#) summarizes the key advantages and limitations of the learning and observational bias approaches discussed in this part of the review.

The learning and observational biases discussed above enable models to accurately capture system dynamics and degradation patterns in a way that is consistent with physical laws. Yet, effective PHM ultimately requires translating these *passive predictions* into *active decision-making* that interacts with and impacts the physical system itself. This transition is achieved through **reinforcement learning** (see [Section 4](#)), which closes the loop by allowing agents to move from model-based predictions back to the physical world. By building on physics-informed

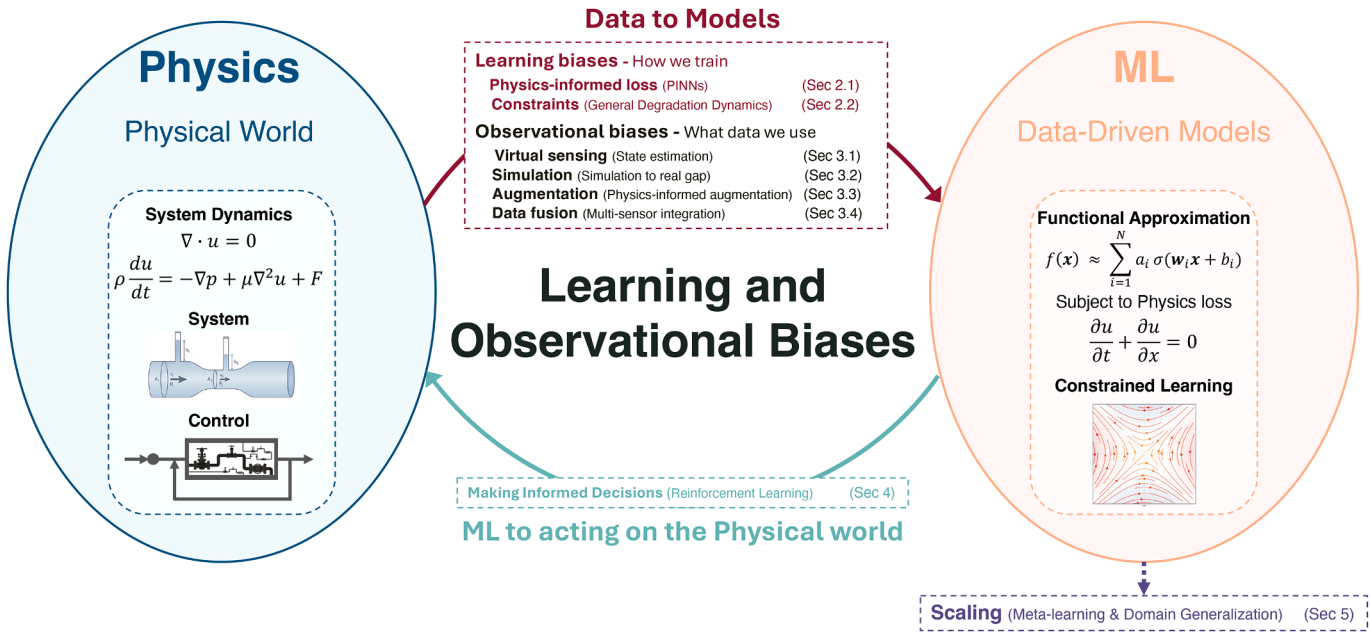


Fig. 2. Overview of how learning and observational biases bridging Physics and Machine Learning within PHM framework. Observational biases (Section 3) ensure data reflects physical system realities, improving model generalization and robustness. Learning biases (Section 2) guide model optimization to produce physically consistent prediction. Reinforcement learning (RL) (Section 4) leverages these biases to transition predictive insights into actionable maintenance decisions, closing the loop between modeling and operational effectiveness in PHM.

Table 2  
Key advantages and limitations of learning and observational biases in PHM (Part II).

Bias	Approach	Key Advantages	Main Limitations
Learning	Physics-Informed (PINNs)	NN Improves generalizability via hard physics constraints; enforces physical consistency; reduces labeled data needs; extrapolation beyond observed regimes.	High computational cost; hyperparameter sensitivity; scalability issues in high-dimensional PDEs; struggles with discontinuities.
	General Degradation	Flexible embedding of general physical properties (monotonicity, trendability); applicable across diverse systems; less restrictive than PDEs.	May oversimplify complex physics; requires domain knowledge for constraint selection; potential for over-regularization.
Observational	Virtual Sensing	Cost-effective alternative to physical sensors; estimates unmeasurable variables; provides redundancy for unreliable sensors; deployable in real time.	Dependent on model fidelity and data quality; domain drift and calibration needs; limited transfer across operating regimes.
	Sim-to-Real	Reduces reliance on labeled field data; covers rare/unsafe failure scenarios; enables safe and controllable data generation.	Simulator fidelity limits transfer; domain gap/negative transfer risks; complex validation.
	Generative Modeling	Generates rare-failure samples and augmentations; improves robustness and class balance; models data manifolds; adaptable across assets.	Computationally intensive; synthetic-real domain gap; propagates dataset biases; training instabilities (e.g., mode collapse).
	Data Fusion	Combines complementary modalities for richer health representation; improves accuracy and robustness; resilience to single-sensor failures.	Heterogeneous data alignment / synchronization is challenging; sensitivity to missing modalities; calibration and scaling issues.

representations, RL agents can learn maintenance and operational policies that not only balance operational constraints, costs, and reliability, but also ensure that decisions are informed by the true physical behavior and degradation processes of the system.

Finally, moving from individual assets to fleet-wide deployment introduces significant **scaling challenges** (Section 5). To address these, we highlight fast adaptation approaches such as meta-learning and few-shot learning (Section 5.1), which facilitate rapid adaptation to new systems with limited data. In addition, domain generalization methods (Section 5.2) are reviewed for their role in ensuring models maintain robust performance across diverse assets and operational contexts. These methods address the practical requirements of large-scale PHM, supporting reliable performance across multiple assets and varied operating environments.

Fig. 2 illustrates our framework, which moves from physics to machine learning and then back to the physical system. Learning and observational biases form the basis for physics-informed predictions, while reinforcement learning closes the loop by translating these predictions into actions that directly affect system behavior. The following sections provide a detailed review of observational and learning biases, including their foundational principles, practical applications in PHM, current limitations, and future research directions. Integrating these systematic biases supports the development of robust, reliable, and scalable PHM solutions, ultimately enhancing asset reliability and enabling optimal operational decisions.

**Main Takeaways.** This review shows that PHM is increasingly evolving from conventional data-driven prediction toward physics-informed, decision-oriented, and fleet-level health management. The

**Table 3**  
Checklist of Physics-informed neural networks (PINNs).

Applications in PHM	Type of tasks	Fault detection; RUL prediction.
	Addressed challenges	Sparse data; Physical inconsistency; Poor generalization ability.
Requirements	Prior knowledge	Governing differential equations.
	Data	Labeled data.
	Type of bias	Learning bias integrated through the loss function.
	Assumptions	Smoothness.
Summary	Advantages	Improved generalizability; Incorporation of physical laws; Reduced dependency on large datasets.
	Disadvantages	Computational cost; Hyperparameter sensitivity; Scalability to high-dimensional problems;
Resources	Repositories	DeepXDE [8], NeuralPDE[9]
	Datasets	PINNacle [10]; PDEBench[11];

field has made significant progress by incorporating learning and observational biases that embed physical knowledge into model training and data representation, enabling more robust and physically consistent predictions under sparse, noisy, or partially observable data. Reinforcement learning has emerged as a key mechanism for translating physics-informed predictions into intelligent maintenance and operational decisions, closing the loop between modeling and physical system behavior. Looking ahead, the integration of physics-informed PHM with Agentic AI systems, capable of autonomous monitoring, reasoning, and decision-making, offers additional opportunities for scalable and adaptive health management. Despite this progress, the field still lacks systematic benchmarking in realistic industrial settings, reliable handling of the simulation-to-real gap, and scalable methods that disentangle degradation from operational dynamics. Building on these gaps, promising research directions include adaptive physics-informed training, more realistic physics-aware generative modeling, hybrid approaches that integrate partial physics with data-driven inference, and scalable generalization across diverse assets. These developments are particularly relevant for safety-critical and physics-governed domains such as energy systems, transportation, civil and subsurface infrastructure, and complex manufacturing systems.

The remaining of this paper is organized into five main sections. Section 2 details leveraging learning biases to constrain model training while Section 3 explores leveraging observational biases through virtual sensing and data augmentation. Subsequently, Section 4 discusses how to close the loop using reinforcement learning for decision-making. Finally, Section 5 addresses scaling beyond single systems via meta-learning and domain generalization, followed by the conclusion in Section 6.

## 2. Leveraging learning bias: Loss functions and constraints for robust PHM models

### 2.1. Physics-informed neural networks (PINNs)

#### Introduction to PINNs

Traditional neural networks (NNs) have become widely used in PHM applications due to their ability to capture complex patterns from data. However, high-quality labeled datasets are often scarce, as collecting real-world measurements or running detailed numerical simulations for complex systems is both costly and time-consuming. In these “low-data” settings, purely data-driven models face key challenges: First, NNs trained on limited data may struggle to generalize or converge at all, leading to unreliable predictions under new or varying conditions. Second, standard NNs can produce physically implausible outputs that violate essential system constraints, such as symmetries and conservation laws, undermining their reliability in PHM applications.

PINNs, first introduced by Raissi et al. [12], address these challenges by embedding physical priors into the training process. These priors,

usually given in the form of governing partial differential equations (PDEs), are encoded as regularizing terms in a model’s loss function. They then act in conjunction with inductive biases to produce predictions that are consistent with established physical principles, rather than relying exclusively on data-driven correlations. This makes PINNs a prototypical example of learning bias, enhancing generalization in low-data regimes.

PINNs were first conceived as coordinate-based Multilayer Perceptrons (MLPs), used as universal approximators for solutions to PDEs. The key idea is to use these networks to directly model solutions, supplementing scarce data with domain knowledge by embedding differential equations into their training loss as residues. For example, given a general differential operator  $\mathcal{L}$  and a source term  $g$ , the physics-based constraint  $\mathcal{L}(u) = g$  can be encoded as a *physics-informed* loss:

$$\begin{aligned} \mathcal{L}(u) &= g \\ \implies L_{PDE} &= \|\mathcal{L}(u) - g\|_2^2 = 0. \end{aligned} \quad (1)$$

which is minimized when the network’s prediction  $u$  satisfies the underlying physical law. Modern machine learning frameworks enable efficient computation of derivatives with automatic differentiation, enabling the evaluation of such loss terms at randomly sampled collocation points throughout the domain. As illustrated in Fig. 3, PINNs combine the physics-based loss with data-driven terms (e.g., initial / boundary conditions or measurement data):


$$L_{PINN} = w_{\text{data}} L_{\text{data}} + w_{\text{PDE}} L_{\text{PDE}} \quad (2)$$

where:

- $L_{\text{data}}$  represents the **data-driven loss**, which measures the discrepancy between the model’s predictions and actual observations, as well as boundary and initial conditions.
- $L_{\text{PDE}}$  enforces the **governing PDE constraints**, ensuring that the neural network satisfies the underlying physical equations at randomly sampled collocation points.
- $w_{\text{data}}, w_{\text{PDE}}$  are weighting factors that balance the contributions of each loss component.

As PINNs were first devised in the context of MLPs for PDEs, these represent the majority of the resources and applications available in the literature (see Table 3). Code repositories for constructing and training MLP-PINNs are available in Pytorch [8] and Julia [9]; likewise, current benchmarks largely focus on the solution of time-dependent PDEs, such as the wave, heat, or Navier-Stokes Equations [10,11]. These, in turn, may be used may be transferred to PHM applications

PINNs have since been considered in a broader sense, inspiring a number of extensions, including schemes for domain decomposition [13] and Bayesian inference [14]. Physics-informed constraints have been integrated as regularizers in a large number of modern architectures, including transformers [15], and Neural ODEs [16]. In this more



**Fig. 3.** Overview of a physics-informed neural network. PINNs are trained to directly approximate solutions to PDEs by encoding them as residues in physics-informed loss terms.

general context, physics-informed terms may take the form of general operators, conservation laws, and symmetries, both continuous and discrete. These terms may encode a variety of physical priors, such as Hamiltonians [17] and continuous symmetries [18]. Integrating such physical insight directly into the training process reduces reliance on large labeled datasets, improves generalization, and ensures that model predictions remain consistent with known physical laws.

#### Application of PINNs in PHM

PINNs and their variants have seen a surge in interest in recent years, leading to a large body of work on PHM, including applications in detection, diagnostics and prognostics, where they contribute to improving the reliability of data-driven frameworks. While diverse, these applications are often based on combining observational data with a wide range of physics-based priors, such as conservation of energy. The concept of physics-informed (as well as "physics-inspired" and "physics-guided") networks has broadened in scope, leading to a diverse range of additional techniques for embedding physical insight into PHM model training, including physics-informed data augmentation [19].

In Sun et al. [20], a physics-informed term enforcing energy conservation is used to enable robust damage detection in wind fan turbines, using this regularization to refine unreliable measurements from acoustic sensors. In [21], PINNs are used to perform non-destructive testing by using the surface acoustic wave equation alongside ultrasound data to detect and describe surface cracks. In [15], PINNs are used to track thermal insulation aging in power-plant transformers, using a PDE-based heat propagation model to extract detailed spatial thermal profiles from localized measurements.

PINNs have also been used for multiple prognostics tasks. For instance, Zheng et al. [22] applied PINNs to predict crack growth in quasi-brittle materials in an unsupervised context. In this case, the loss function was formulated in terms of differential equations modelling damage increase and the minimization of elastic potential energy, leading to physically plausible crack configurations. Similarly, Wen et al. [23] developed a model fusion approach using PINNs for RUL predictions of Lithium-ion batteries. Here, the physics-informed loss encodes coarse low-dimensional approximations for degradation dynamics, while a data-driven model is used to model spatial dependences across the power cell. Additionally, Wang et al. [24] introduced a PINN model for state-of-health estimation, whereby both the physics-informed equation residue as well as its solution are learned with neural networks, leading to a flexible framework for learning the unknown dynamics of battery degradation over time. A summary of the application contexts for PINNs is provided in Table 3.

#### Advantages and limitations of PINNs

These applications highlight the unique advantages of PINNs in PHM. By blending physics-based constraints with data-driven learning, PINNs enable more reliable and interpretable predictions, particularly in scenarios where high-quality labeled data are limited, costly, or noisy.

Their ability to enforce physical consistency during training means that PINNs naturally denoise and correct sensor measurements, leading to robust predictions that remain consistent with the underlying system dynamics. Unlike purely data-driven models, which often require large datasets to capture degradation patterns, PINNs leverage governing equations to guide learning. This allows them to predict fault onset, degradation progression, or remaining useful life accurately even from limited or incomplete data. Nevertheless, while they offer significant advantages in PHM, they also have several limitations that need to be addressed for broader adoption.

A significant limitation of PINNs is their sensitivity to hyperparameters and the challenge of loss balancing [25]. Because PINNs combine data-driven and physics-based losses, the choice of weighting factors is critical: poorly tuned weighting factors can cause instability during training or lead to suboptimal solutions. If the physics loss is weighted too heavily, the network may ignore available data; conversely, if the data loss dominates, the model risks losing physical consistency and behaving like a purely data-driven network. Another challenge is the requirement for explicit, accurate governing equations. In many practical PHM applications, such as material fatigue or degradation prediction, these equations may be incomplete, approximate, or entirely unknown, limiting the effectiveness of PINNs. Furthermore, relying solely on physics-informed regularization is often insufficient to enforce physically plausible solutions. In low-data regimes, PINNs may fail to converge or produce trivial, non-physical outputs. This is particularly prevalent for systems dealing with high-frequency behavior or complex geometries, restricting PINNs to relatively simple equations. This often requires the introduction of additional inductive biases, such as specialized neural architectures [26] or tailored regularization strategies.

Despite these limitations, ongoing research in adaptive loss balancing, hybrid modeling approaches, and more computationally efficient training algorithms continues to broaden the applicability and scalability of PINNs in PHM.

#### Unexplored or emerging applications of PINNs in PHM

Although Physics-Informed Neural Networks have been widely used across multiple sub-domains of PHM, they still hold considerable untapped potential, particularly in domains where classical degradation models exist but are computationally expensive, incomplete, or difficult to integrate with real-world data.

So far, in part due to the limitations of PINNs, the physics-informed losses used often refer to relatively basic, phenomenological descriptions of the system. For this reason, they still cannot be considered replacements for traditional numerical simulators when complex geometries or dynamics are at play. Likewise, the application of physics-informed terms to the training of emerging specialized architectures remains a promising field in general.

In terms of applications, the analysis of tribological wear prediction in gears and bearings shows promise. Traditional wear models, such as Archard's law, provide general insights but fail to capture complex tribological interactions like nonlinear friction and lubrication effects [27]. PINNs have been used in an unsupervised manner in this context [28] and could be used further to integrate physical laws with sensor data to improve the real-time estimation of wear rates.

In corrosion and material degradation modeling, PINNs could enhance the understanding of time-dependent electrochemical reactions affecting structures like pipelines, aircraft, and marine vessels. Modeling corrosion-fatigue has proven to be hard, which has led to the development of hybrid data-driven frameworks [29]. PINNs could be used in

conjunction with real-world data to account for environmental variability (e.g., humidity, salinity) and provide more robust and site-specific degradation forecasts.

Finally, PINNs could be used to integrate a wider range of physical priors with non-destructive testing (NDT). NDT often involves the interaction of multiple complex and interdependent physical phenomena, which can be described by elaborate PDE systems. PINNs could improve the detection of hidden subsurface defects using ultrasonic, eddy current, or infrared thermography data, along with the mathematical descriptions thereof, leading to more reliable condition assessments and defect identification in aerospace, civil, and industrial structures.

#### Promising research direction for PHM: Physics-informed computer vision

A particularly promising direction for extending PINNs in PHM is Physics-Informed Computer Vision (PICV). Visual data is increasingly central to PHM applications, particularly for inspection tasks such as defect detection [30], structural monitoring [31,32], and visual condition assessment [33]. With the growing use of autonomous platforms (e.g., drones, underwater vehicles, robotic arms), computer vision (CV) enables scalable, non-intrusive monitoring across diverse assets and environments.

Despite this potential, current CV models in PHM are often limited by their reliance on large labeled datasets and poor generalization across operating conditions. Many models are trained for specific scenarios and may underperform under variable viewpoints, lighting, or noise—frequent challenges in real-world deployments. PICV aims to address these limitations by embedding physical constraints into the learning process, improving robustness, and interpretability in data-scarce or safety-critical settings. This integration of physics with vision offers a natural extension of PINNs to visual modalities, opening new opportunities for physics-consistent PHM systems. In practice, PICV can enable PHM models to move beyond passive visual inspection toward estimation of stresses, deformations, flows, or thermal behavior—quantities directly relevant for diagnostics, prognostics, and RUL prediction.

PICV integrates domain knowledge into data-driven CV models by embedding physical priors through three main types of biases [5,34]: (1) Observational biases: capturing physical quantities directly through enriched sensing modalities (e.g., depth or thermal imaging [35,36]); (2) Learning biases: regularizing training using physical equations (e.g., heat transfer, Navier-Stokes) or geometric constraints (e.g., epipolar geometry, photometric consistency) to improve data efficiency and enforce physically plausible outputs [37,38]; (3) Inductive biases: embedding symmetries and invariances (e.g., rotational, scale) into architectures using group convolutions or steerable CNNs [39]. These strategies typically enhance model robustness, learning efficiency, and improve physical consistency, particularly beneficial for PHM tasks where deployments face highly variable conditions. In label-scarce scenarios, common in PHM, these physics-based objectives can also be applied in a self-supervised manner, allowing models to learn from abundant unlabeled visual data while maintaining physical consistency. For an extensive survey on PICV, we guide the reader to [34].

While CV is widely used in PHM across domains—such as railway inspection [40], aerospace surface analysis [30], and civil infrastructure monitoring [31,32]—the explicit use of PICV for PHM remains limited. A few hybrid approaches exist, such as combining grease condition analysis via vision with physics-based models for bearing degradation in wind turbines [41,42], but these are still rare.

PICV opens promising research directions in PHM. Potential applications include dynamic process monitoring, where vision-based fluid flow or thermal field modeling could enable real-time anomaly detection (e.g., cavitation, heat build-up) [37,43]. Additionally, equivariant architectures (e.g., to rotation and scale) can improve inspection robustness for structures and components [33]. Finally, combining vision with sensor data (e.g., vibration, pressure)—a strategy known as multi-modal fusion (see Section 4)—can lead to more comprehensive health

**Table 4**  
Checklist of Degradation dynamics-informed methods.

Applications in PHM	Type of tasks	Health indicator (HI) construction; RUL prediction.
	Addressed challenges	Lack of direct degradation labels; Truncated or noisy trajectories; Weak degradation signals in raw sensor data.
Requirements	Prior knowledge	Fundamental degradation properties (Monotonicity, Trendability, Prognostability, Robustness).
	Data	Time-series data
	Type of bias	Learning bias incorporated through auxiliary loss functions.
	Assumptions	Degradation evolves irreversibly and correlates with accumulated usage.
Summary	Advantages	Does not require explicit PDEs; Effective in data-scarce settings; Produces interpretable HIs; Robust to sensor noise.
	Disadvantages	Requires careful hyperparameter tuning.
Resources	Repositories	<b>HConstruct_optimize</b> [46] (constraint-based health indicator construction), <b>UnsupervisedHI</b> [47] (unsupervised degradation-consistent health indicator learning), <b>PINN4SOH</b> [48] (physics-informed battery state-of-health modeling), <b>Wang2024PINN_RUL_Battery</b> [24] (battery degradation dataset and benchmarking)
	Datasets	XJTU-SY bearing run-to-failure dataset [49]; XJTU-Battery Dataset [50] NASA C-MAPSS [51] and N-CMAPSS turbofan degradation datasets[52];

assessment systems [44,45]. Practically, PICV can facilitate visual wear tracking, crack-growth modeling, corrosion progression estimation, and the monitoring of other degradation processes where physics drives the system's evolution and its appearance.

Despite these opportunities, PICV faces practical challenges, many of them similar to those faced by PINNs. Designing physics-consistent architectures requires domain knowledge in physics (e.g., PDEs, system symmetries), which may be complex to formalize. Some PICV approaches require explicit governing equations, which may be hard to derive or too simplified for real-world PHM. Vision-specific issues also arise: physical constraints must be enforced in image space, where camera projection effects, occlusions, and lighting variations complicate modeling. Finally, PICV methods often require synchronizing visual and physical data across modalities, adding complexity to training and deployment.

## 2.2. Degradation dynamics informed prognostics methods

### Introduction to degradation dynamics-informed prognostics methods

Industrial and infrastructure assets typically undergo progressive degradation over time due to wear, fatigue, thermal stress, and other environmental and operational factors. Generally, such degradation can be modeled using partial differential equations or other first-principles approaches. For example, in the case of wear, PDEs may capture erosion depth as a function of applied loads and microstructural properties [53]. Similarly, battery degradation can be modeled through electrochemical reaction kinetics [54], and turbine blade erosion can be described by stress-strain relationships [55]. However, for complex systems comprising multiple interacting components and difficult-to-measure

parameters (such as friction coefficients, stiffness, damping, contact forces, and material inhomogeneities) PDE-based degradation models become impractical [56]. This is due to the high dimensionality arising from the need to solve coupled PDEs across different components, the strong nonlinear interactions between these equations, and the limited accessibility of internal variables such as internal loads and stresses that are required for model calibration [53].

Instead of directly modeling degradation through first principles, many studies rely on *health indicators*, quantities that aggregate the level of degradation based on sensor data [57,58]. These may be physically interpretable (e.g., crack length, battery capacity) or latent variables learned from multi-sensor signals. The learning process depends on the nature of available data: when direct measurements of degradation are available, health indicators can be learned in a supervised manner [59]; otherwise, semi-supervised or unsupervised techniques are employed to uncover underlying degradation trends from unlabeled data [60–62].

However, engineering systems often operate under data-scarce conditions due to limited sensor coverage or infrequent monitoring, making it difficult to learn reliable health indicators purely from observed sensor data. In such scenarios, incorporating prior knowledge about degradation dynamics becomes essential [63]. This knowledge is grounded in empirical observations that degradation under cyclic loading often follows predictable patterns (such as steady wear, increasing crack length, or decreasing battery capacity) that reflect underlying physical mechanisms [57,63,64].

These patterns are characterized by key properties commonly observed in engineered systems [57,65]: **Monotonicity**, where damage progresses irreversibly (e.g., cracks do not spontaneously heal); **Trendability**, where degradation correlates consistently with operational variables like time or usage cycles; **Prognosability**, where systems of the same type exhibit comparable degradation trajectories under similar conditions; and **Robustness**, where health indicators remain stable in the presence of noise or environmental variation. These properties can be mathematically formalized and embedded as learning biases in the training objective, guiding the learning of physically meaningful health indicators in scarce data settings. Fig. 4 illustrates this constraint-guided learning process, showing how auxiliary loss terms progressively transform non-physical health indicator behaviors into representations that satisfy key degradation dynamics. Below, we describe how these degradation patterns can be incorporated as auxiliary losses to guide the learning of physically meaningful health indicators:

**Monotonicity.** Irreversible degradation processes—such as crack growth or battery capacity decline—can be modeled by enforcing a monotonic change over time in the health indicator. This can be imposed using a penalty term that discourages non-monotonic transitions [57,60]:

$$L_{\text{mono}} = \sum_{t=1}^{N-1} \max(0, H_t - H_{t+1}) \quad (3)$$

where  $H_t$  is the health indicator at time  $t$ . This loss, designed for decreasing health indicators, penalizes instances where the indicator increases (i.e., when  $H_t > H_{t+1}$ ) thereby discouraging non-monotonic behavior.

**Trendability.** Degradation processes typically evolve consistently over time or with accumulated usage. For instance, bearing degradation should correlate with the number of revolutions, and battery degradation should align with charge-discharge cycles. To reflect this, health indicator can be enforced to correlate with the operational cycle index  $C$  [60,66]:

$$L_{\text{trend}} = -\rho(C, H) \quad (4)$$

where  $\rho(C, H)$  measures the degree of correlation.

**Prognosability.** Systems undergoing the same degradation process should exhibit comparable health trajectories that enable consistent prognostic performance. When the health indicator is scaled between 0 (healthy) and 1 (failure), the health indicators should converge to similar values at key degradation stages, particularly at failure. To enforce this, a penalty can be applied to align the indicator values at the start and end of the degradation trajectory [67,68]:

$$L_{\text{consistent}} = (H_{\text{start}} - 1)^2 + (H_{\text{end}} - 0)^2. \quad (5)$$

**Robustness.** Degradation patterns are typically robust to noise and fluctuations in sensor readings caused by environmental or operational variability. To reflect this, the learned health indicator can be constrained to remain smooth over time, suppressing high-frequency variations that do not correspond to actual degradation. This behavior can be encouraged using regularization or temporal smoothing techniques [68,69]:

$$L_{\text{robustness}} = \frac{1}{N} \sum_{t=1}^N \exp\left(-\left|\frac{H_t - \bar{H}}{\bar{H}}\right|\right) \quad (6)$$

where  $\bar{H}$  is the mean health indicator over time. For modeling a system's degradation process, one or more of these loss formulations can be selectively applied depending on the available data, domain knowledge, and desired properties of the health indicator. Table 5 summarizes these concepts, providing physical interpretations and visual intuitions for each mathematical property.

#### Application of degradation-dynamics-informed methods in PHM

Recent studies have incorporated degradation dynamics into the learning process by embedding monotonicity constraints in the health indicator formulation along with supervision on the underlying measured parameters. For example, Kim et al. [70] used simulated crack growth data based on Paris' law to constrain a neural network to follow the irreversible progression of damage, resulting in strong agreement with known crack propagation trends. Building on similar physics-motivated constraints, Chen et al. [71] developed a degradation-consistent recurrent neural network for bearing prognosis using temperature signals. By imposing a monotonic trend on predicted temperature values, their method improved the temporal consistency of degradation modeling and yielded more stable remaining useful life estimates.

Other studies have employed unsupervised or semi-supervised approaches to learn latent health indicators from multi-sensor data by imposing degradation-informed constraints, particularly in settings with no direct measurement of degradation. For example, Wang et al. [72] proposed a deep learning-based sensor fusion method that integrates signals such as vibration and acoustic emissions, applying monotonicity and range constraints to produce a consistent indicator of degradation. Their approach was demonstrated on simulated turbofan engine degradation data. Similarly, Chen et al. [68] developed a feature fusion framework for lithium-ion batteries and rolling element bearings, optimizing the combination of degradation-relevant features while enforcing monotonicity, trendability, and robustness. Following a similar idea, Bajarunas et al. [66] introduced degradation-informed learning biases (such as trendability and negative gradient constraints) into an autoencoder architecture to infer latent degradation signals from turbofan and battery datasets. Further, Wen et al. [60] constructed a signal-based indicator for aircraft turbine engines by optimizing for trajectory consistency and enhancing signal range across degradation cases. More recently, Qin et al. [73] proposed a supervised multi-head self-attention autoencoder for machinery prognostics, using power function-inspired constraints to improve similarity-based remaining useful life prediction for wind turbine gearboxes and aero engines. These studies demonstrate the utility of degradation-consistent learning objectives for inferring health indicators in complex systems where direct measurements are limited or unavailable. Table 4 summarizes the applications, advantages, disadvantages, and representative resources of

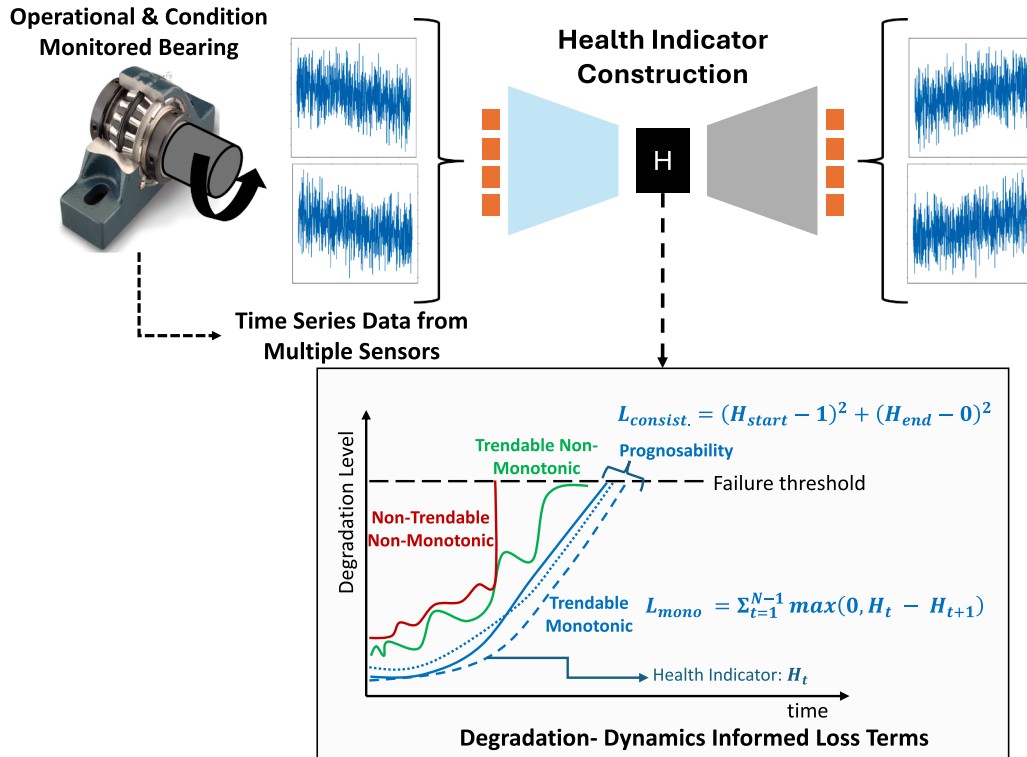


Fig. 4. Overview of degradation dynamics-informed prognostic methods. Time-series sensor data from monitored systems (e.g., bearings) are processed to construct health indicators that progressively satisfy key degradation properties: monotonicity (irreversible decline), trendability (correlation with operational cycles), and Prognosability (comparable endpoints across similar units). Auxiliary loss terms enforce these physics-motivated constraints during the learning process, guiding the evolution from non-physical to physically meaningful health indicators.

Table 5  
Interpretation of degradation-dynamics loss functions.

Property	Mathematical Concept	Physical Interpretation	Visual Concept
<b>Monotonicity</b>	$\sum_t \max(0, H_t - H_{t+1})$	Degradation is irreversible; health should not recover spontaneously.	
<b>Trendability</b>	$-\rho(C, H)$	Degradation correlates with usage; health follows a consistent trend.	
<b>Prognosability</b>	$(H_{end} - 0)^2$	Units of the same type fail at comparable degradation levels.	
<b>Robustness</b>	Smoothness penalty on $H_t$	True degradation evolves slowly; spikes are likely sensor noise.	

degradation-dynamics-informed prognostic methods. The listed repositories provide open-source implementations of constraint-based health indicator construction and physics-informed degradation modeling, including optimized health indicator fusion with degradation constraints (HIconstruct\_optimize [46]), unsupervised health indicator learning guided by degradation consistency (UnsupervisedHI [47]), and physics-informed neural networks for lithium-ion battery state-of-health estimation (PINN4SOH [48], Wang2024PINN\_RUL\_Battery proposed in [24]). The datasets comprise widely used run-to-failure benchmarks spanning multiple domains, including the XJTU-SY bearing degradation dataset [49], the XJTU open-access lithium-ion battery dataset [50], and the NASA C-MAPSS [51] and N-CMAPSS [52] turbofan engine degradation datasets, which support the evaluation of degradation-consistent health indicator learning and prognostic performance.

#### Advantages and limitations of degradation-dynamics-informed methods in PHM

Degradation-dynamics-informed methods incorporate generic constraint into the learning objective to better capture the intrinsic progression of degradation. These inductive biases are particularly beneficial in high-dimensional, multi-sensor settings, where each sensor provides only partial and potentially noisy observations of the underlying degradation process that is typically not directly measured or known [58,66,69]. By embedding these constraints, such methods improve robustness to sensor noise, mitigate overfitting, and enhance the interpretability of the learned health indicator.

Owing to the generic nature of these learning biases [69], degradation-dynamics-informed methods are broadly applicable across diverse domains-ranging from bearings [74] and batteries [75] to turbine engines [66]-thus providing another key advantage for practical deployment in PHM.

However, these methods also present notable challenges. A key limitation lies in their reliance on the assumption of monotonic degradation, which does not universally hold. Systems subject to resets (e.g., maintenance events), regime changes, or operational variability may exhibit non-monotonic degradation trajectories [76–78]. In particular, systems exposed to random external shocks or varying operational loads deviate from smooth, unidirectional degradation patterns, thereby violating the foundational assumptions of these methods. Moreover, the varying operating conditions can make it difficult to consistently impose these constraints across different operational regimes, potentially leading to inaccurate prognostic outcomes when the assumed degradation dynamics do not align with the actual system behavior.

#### Unexplored or emerging applications of degradation-dynamics-informed methods in PHM

Degradation-dynamics-informed prognostic methods have shown strong performance in established applications such as bearings [74], batteries [75], and turbofan engines [66]. However, their application to real industrial data collected under realistic operating conditions remains limited, with many studies relying on controlled experimental setups or simulated datasets (e.g., NASA's turbofan simulation data [51]). Moreover, their use in multi-component systems and more complex variable industrial environments is underexplored. Emerging domains such as wind turbines, railway components (e.g., axle bearings and brake pads) [79], and industrial robotic systems [80] operate under non-stationary conditions involving varying loads, control strategies, and maintenance events. These factors often lead to non-monotonic degradation patterns that violate assumptions underlying many current methods.

To address these challenges, future research should focus on combining *regime-aware segmentation* with *partial-monotonic constraints*, or explicitly incorporating control inputs into the learning architecture.

**Table 6**  
Checklist of *Soft sensing*.

Applications in PHM	Type of tasks	Fault detection and Prognostics.
	Addressed challenges	Lack of direct measurements; Sensor costs; Sensor failure or degradation; Estimation of difficult-to-measure or unmeasurable variables; Data sparsity or missing values.
Requirements	Prior knowledge	Optional for Data-Driven: System dynamics; Physics-based models.
	Data	Labeled data; Historical sensor measurements.
	Type of bias	Observational bias.
	Assumptions	System behavior can be approximated through mathematical or data-driven models.
Summary	Advantages	Cost-effective alternative to physical sensors; Can estimate unmeasurable variables; Backbones for unreliable physical sensors.
	Disadvantages	Sensitivity to model inaccuracies in Model-Driven methods, Dependence on data quality; Potential interpretability issues in deep learning models.
Resources	Repositories	GgNet [82], PEGNN [83], HTGNN [84],
	Datasets	Train-Bridge-Track Interaction [84], Bearing Load Prediction [84], District Heating Networks [83,85]

By identifying operational segments where degradation evolves consistently (using tools such as change point detection, load profile analysis, or maintenance records) constraints like monotonicity or trendability can be selectively enforced. Prior work has highlighted the importance of jointly modeling gradual degradation and shock-induced damage, particularly in multi-component systems where failures may result from both cumulative wear and abrupt events [76,81].

Integrating such physics-informed and shock-aware mechanisms into deep learning frameworks-as auxiliary loss functions-can enhance the robustness, interpretability, and adaptability of learned health indicators. These advances are essential for deploying degradation-dynamics-informed methods in realistic industrial environments characterized by variable operating conditions, irregular maintenance, and noisy sensor data.

### 3. Leveraging observational biases: Augmentation and physical constraints

#### 3.1. *Soft sensing*

##### *Introduction to soft sensing*

In industrial settings, sensors are widely used to monitor processes, assess health condition, and ensure asset reliability. Ideally, placing sensors at every critical location would provide the most comprehensive process insight. However, practical and economic constraints often make this impossible. Some areas are inaccessible or exposed to harsh environments, where installing or maintaining sensors is not feasible. Additionally, retrofitting existing systems to accommodate new sensors can be prohibitively expensive and disruptive, often requiring production downtime, equipment modifications, or significant re-engineering. As a result, many important system variables cannot be measured directly by physical instruments alone. To address these challenges, *soft sensing* (also known as *virtual sensing*) has emerged as a practical alternative. *Soft sensing* uses computational and domain knowledge to

estimate key process variables or product quality based on available measurements [86], as illustrated in Fig. 5. This approach offers a cost-effective solution when physical sensing is impractical or prohibitively expensive [87]. Selecting the right soft sensor model is crucial for ensuring accurate predictions and effective system monitoring. The two main approaches for developing soft sensors are white-box (model-driven) and black-box (data-driven) modeling, each with specific advantages and limitations [88].

**Model-driven, or white-box, soft sensing approaches** are grounded in fundamental principles of physics, chemistry, and engineering that govern system behavior [89]. Common techniques include first-principle models, Kalman filters, and observers [86]. **First-principle models** are particularly effective when the governing equations of a system are well-understood and can be mathematically formulated, providing high interpretability through conservation laws [90]. Their main limitation, however, is the need for comprehensive system knowledge. **Kalman filters** provide a complementary approach when a complete mechanistic description of the system is not feasible (see Section 4 of Part 1). These estimation algorithms fuse noisy sensor measurements with simplified system models, which may be based on linearized versions of nonlinear dynamics, mass or energy balances under steady-state assumptions, or empirical relationships derived from experimental data. By capturing the essential system behavior without requiring a full mechanistic model, Kalman filters infer unmeasured states through sequential estimation (see Section 4), continuously updating their predictions as new measurements arrive to enable real-time state tracking [91,92].

**Data-driven approaches** construct soft sensing models directly from empirical data, without relying on explicit physical models of the underlying processes [93]. They are particularly valuable in industrial contexts where detailed process models are unavailable or incomplete, but large volumes of operational data are routinely collected.

Historically, linear and piecewise linear models have been the foundation of data-driven soft sensor modeling [94,95]. Among these, Partial Least Squares (PLS) is a widely used method in industrial practice due to its robustness to measurement noise and multicollinearity [96]. However, the linear assumptions of PLS limit its capacity to represent complex, nonlinear system dynamics. To address this, nonlinear approaches such as Support Vector Regression have been explored [97–99], along with probabilistic approaches like Gaussian Process Regression, which can capture nonlinear relationships while also providing uncertainty estimates [100,101]. Nonetheless, these machine learning methods typically require extensive feature engineering, which often falls short of fully capturing the complexity of underlying process behaviors.

Recently, deep learning-based techniques have gained traction for soft sensing, as they can automatically learn relevant features and representations from raw data. Unlike conventional methods, deep learning reduces reliance on expert-crafted features [102,103]. This shift has opened new possibilities for modeling highly complex and nonlinear processes.

To capture spatial correlations in high-dimensional process data, convolutional neural networks (CNNs) have been applied to extract multiscale spatial features, with several studies demonstrating in taking advantage of the correlations between process variables [104,105]. To model temporal correlations, gated recurrent units (GRUs) have been adopted to capture sequential dependencies in process data, with bidirectional and hybrid designs improving predictive accuracy and robustness [106,107]. However, these conventional deep learning models generally assume fixed data structures and may not fully capture the irregular, dynamic, and relational nature of complex sensor networks.

To address these limitations, both transformers and graph neural networks (GNNs), including their extension Spatial-Temporal Graph Neural Networks (STGNNs) (see Section 2 of Part 1 [6]) have been explored. Transformers are well suited for capturing global, long-range dependencies through attention mechanisms [108], whereas GNNs explicitly encode relational inductive biases by leveraging the topology of sensor

networks. STGNNs further extend this capability by jointly modeling spatial structures and temporal dynamics [109]. For example, Felice et al. [82] proposed a graph-based deep learning architecture for soft sensing in sparse multivariate sensor networks. It models both spatial and variable-wise dependencies using a nested graph structure: an inter-location graph captures learned similarities between sensor sites, while an intra-location graph models relationships among sensor variables at each location. While early graph-based soft sensing methods mainly focus on learning structural relationships among sensors, later studies extend these approaches by explicitly incorporating causal reasoning to better capture spatial-temporal dynamics and improve model robustness in industrial processes. In particular, causal graph learning enables the identification of meaningful cause-effect relationships among process variables, providing more stable and interpretable soft sensor models. A representative example is the development of stable soft sensors for batch processes based on causal inference and graph convolutional networks, where both global and local causal relationships are learned and aggregated through causal graph convolutions to reduce intra-batch and inter-batch prediction errors [110]. Building on this causal graph-based paradigm, Hu et al. [111] further extend soft sensing to heterogeneous industrial settings by proposing a Heterogeneous Causal Graph Attention Network for mass customization manufacturing. Their approach integrates causal discovery with multi-scale sensor modeling to accurately predict product quality across customized products with varying assemblies.

Another important advancement in soft sensing is the incorporation of physical laws and domain knowledge directly into the learning process. Niresi et al. [83] propose a framework that augments the graph structure with auxiliary nodes and connections derived from governing equations, embedding process physics into the model architecture. Adding auxiliary nodes to the graph structure allows the GNN to map from a more enriched input space to the soft sensors, which improves estimation accuracy. Moreover, physics-guided graph learning combines a topology informed by known process relationships with learned correlations, and constrains the training objective through physics-based regularization terms that penalize violations of conservation constraints [112], which leads to more accurate soft sensor estimation. Beyond graph-based methods, physically informed hierarchical learning [113] integrates reduced-order physical models into a multi-level neural architecture, using them as inductive biases to structure intermediate representations and guide parameter estimation when the underlying PDE dynamics are only partially known, further enhancing soft sensing performance.

#### *Applications of Soft Sensing in PHM*

In PHM, soft sensing serves multiple critical functions that enhance system reliability and reduce operational costs. Its applications span diverse domains, including process industries, mechanical systems, structural health monitoring, and energy systems, each with unique measurement challenges and data characteristics.

**Estimation of Unmeasurable Variables:** Soft sensors enable the estimation of variables that cannot be directly measured due to harsh environments, physical inaccessibility, or prohibitive costs. In process industry, they are widely used to monitor complex product quality variables such as chemical component concentrations, polymer density, and melt flow rate, which are essential parameters for precise control and optimization [114]. Beyond process environments, structural health monitoring is another key domain for soft sensing. In bridges, limited sensor deployments often constrain direct measurement of variables such as strain, displacement, or real-time load. The number and placement of sensors are often further restricted by costs, deployability, or accessibility, which motivates the use of virtual sensing approaches. Virtual sensors can infer these unmeasured responses by fusing heterogeneous sensing modalities with environmental and loading context. Zhao et al. [84] propose a heterogeneous temporal GNN that models distinct

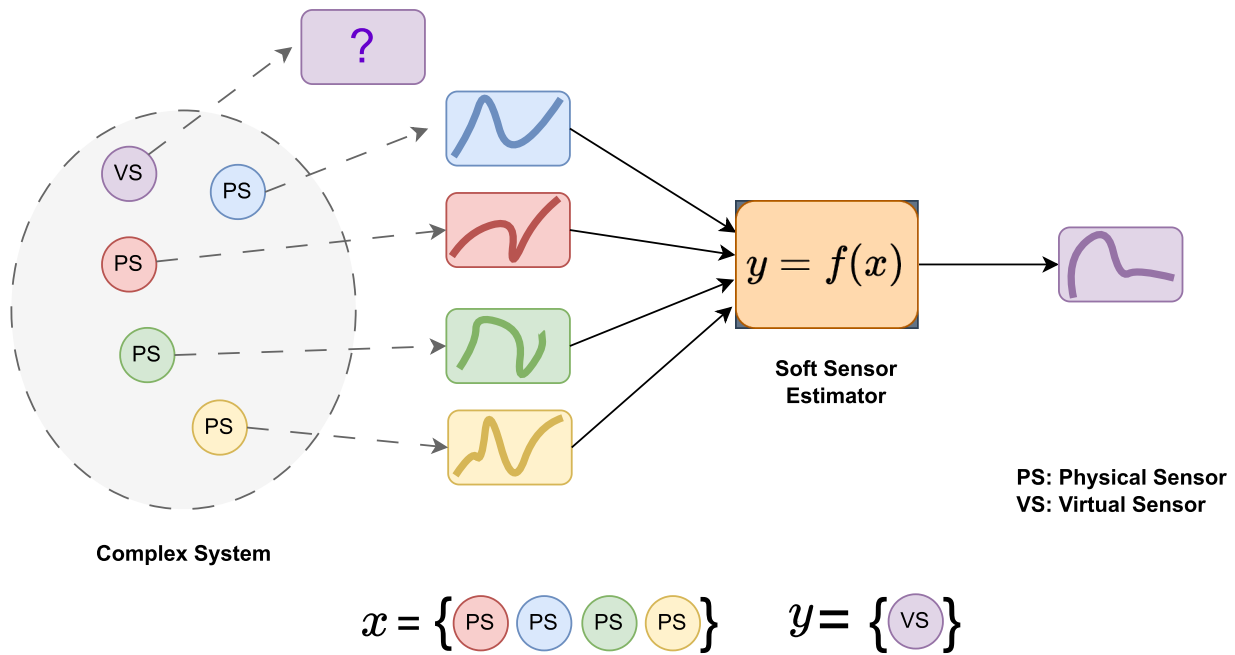


Fig. 5. Illustration of a complex system with both physical sensors (measuring real-time parameters) and virtual (soft) sensors. The diagram highlights the mapping process from physical sensors to soft sensors using deep learning or other predictive modeling techniques, enabling enhanced monitoring, fault detection, and system optimization.

sensor types, integrates operating condition context, and predicts both bearing loads in rotating machinery and real-time loads in bridges, demonstrating the cross-domain applicability of advanced GNN-based virtual sensing. Virtual sensing was also applied to electromechanical assets, with electric motors employing soft sensors to infer internal winding temperatures and insulation aging [115,116]. Moreover, energy storage systems utilize virtual sensing in battery management to derive crucial state-of-charge and state-of-health estimates from operational voltage, current, and temperature signals [117]. These diverse applications highlight the adaptability of soft sensing across domains, enabling the indirect measurement of critical variables that would otherwise remain inaccessible.

**Backup Measurements for Sensor Failure Prevention and Detection:** Soft sensors serve also as backup measurements that maintain operational continuity when physical sensors fail, preventing costly shutdowns. In addition, they support the detection of faulty or degraded sensors.

Darvishi et al. [118] proposed a data-driven modular framework, in which the available sensor signals are partitioned into a reliable set and an unreliable set. Multilayer perceptrons are trained to model each signal in the unreliable set using only the reliable signals as input. By comparing the predicted (virtual) signals to the actual sensor outputs, the system performs residual analysis to detect faults. In a different application, district heating networks rely on differential pressure to monitor pump performance, identify flow imbalances, and detect faults such as blockages or leaks. Accurate measurement of this parameter is essential for optimizing control strategies and maintaining energy efficiency, and data-driven soft sensors are employed to infer differential pressure values [119]. Moreover, in HVAC systems, soft sensors support fault tolerance (i.e., maintaining reliable system outputs despite sensor faults) by detecting sensor anomalies through residual analysis and maintaining reliable outputs using soft sensors that are trained based on historical data [120]. This ensures continued operation and control accuracy despite sensor failures, improving system resilience and reducing maintenance needs.

**Data Augmentation for Downstream Tasks:** Soft sensing can also expand the set of input variables by inferring additional unobserved

quantities that can be used for more accurate and more interpretable fault detection, isolation and prognostics [2]. Arias et al. [4] employ a calibrated physics-based performance model of a turbofan engine to estimate latent health-related parameters and virtual signals, which are then combined with sensor measurements as inputs to a deep learning model. This hybrid approach enhances prognostic performance while reducing dependence on large training datasets.

A summary of the application contexts for soft sensing is provided in Table 6. We include several GNN implementations that are particularly relevant to soft sensing in PHM applications. These implementations encompass heterogeneous GNNs (HTGNNs) and physics-enhanced GNN (PEGNN) developed using PyTorch Geometric [121,122], as well as GgNet, which is implemented using Torch Spatiotemporal [123].

In addition, we summarize a set of publicly available datasets on soft sensing in PHM. The District Heating Network dataset [83,85] is a simulation-based dataset that includes measurements from various sensors, such as mass flow rate, pressure, and temperature sensors, used to monitor the operation of a district heating system. In this dataset, soft sensing tasks are typically formulated to estimate a subset of sensor variables, for example temperature and pressure, from other available measurements, such as mass flow rate sensors.

The Bearing Load dataset [84] is a real-world experimental dataset collected from a bearing test rig operating under diverse load and rotational speed conditions. It includes temperature and vibration sensor data and is commonly used to predict axial and radial bearing loads.

Finally, the Train-Bridge-Track Interaction [84] dataset is a comprehensive, year-long simulated dataset designed for structural health monitoring applications. It contains displacement and acceleration sensor readings collected under varying ambient temperature conditions and is used for bridge live load prediction.

#### Advantages and Limitations of Soft Sensing

Soft sensing has become vital in PHM, where it enables estimation of health indicators and process variables that are difficult or impossible to measure directly. Its main advantages include: (1) **Enhanced observability:** estimation of variables that cannot be directly measured,

including internal degradation states, operating loads that influence degradation, process conditions at additional spatial locations, and other critical unmeasured variables. (2) **Cost and feasibility:** reduced reliance on expensive, intrusive, or failure-prone physical sensors, allowing continuous monitoring in situations where installing or maintaining direct measurement devices would be impractical. (3) **Data completeness and enrichment:** reconstruction of variables already monitored by physical sensors to fill temporal or spatial gaps caused by intermittent or faulty measurements, and generation of derived features (such as health indices) that provide additional, physically meaningful inputs for PHM applications.

Despite their advantages, soft sensing in PHM also face several limitations. Model-driven approaches require accurate system knowledge and reliable calibration data. Developing and maintaining such models for complex assets (e.g., wind turbines, jet engines, or manufacturing equipment) is often labor-intensive and may be impractical when system behavior is nonlinear, uncertain, or poorly characterized. Furthermore, white-box models typically rely on steady-state operations and therefore struggle to adapt to dynamic, real-world PHM scenarios, which limits their predictive accuracy during transients or unexpected operating conditions

While Data-driven methods have become popular in PHM for their ability to capture complex, nonlinear relationships from historical or operational data, their success relies on the availability of large, diverse, and high-quality datasets. In practice, assets often undergo evolving degradation modes and operate under varying conditions, making it difficult for purely data-driven soft sensors to generalize reliably without frequent model updates.

#### Emerging or unexplored applications of soft sensing in PHM

Current research in PHM soft sensing is exploring several promising directions. One of the promising strategies for soft sensing in PHM is using PINNs (see Section 2.1), and their graph-based counterparts PIGNNs (see Part I [6]), which integrate physical laws, such as conservation principles or degradation equations, directly into model architecture. This approach enables soft sensors that remain consistent with known dynamics even in data-scarce environments or under distribution shifts. Foundation models represent another promising direction. They can be prompted, through in-context learning (see Section 5.1), with structured time-series data, process descriptions, and historical logs to perform tasks like auxiliary variable selection, uncertainty-aware prediction. Multimodal (see Section 3.4) foundation models can also combine heterogeneous data sources, using language to support transparent reasoning [124]. These techniques reduce reliance on large labeled datasets by incorporating domain knowledge in natural language form. These advanced soft sensing methods have significant potential in PHM, where robustness, data efficiency, and interpretability are critical. Another promising direction is transfer learning or domain adaptation (see Section 3.2) which allows models developed on one fleet or system to be adapted to others with minimal new data [125].

#### 3.2. Overcoming the simulation-to-real gap

##### Introduction to overcoming the simulation-to-real gap

Physics-based approaches have long been fundamental in PHM, providing interpretable, generalizable, and physically grounded models for fault detection [126,127], fault diagnosis [128,129], and prognostics [130,131]. By leveraging first-principles models (e.g., governing equations) and domain expertise [132,133], these methods often achieve high predictive accuracy with relatively limited data requirements. However, their effectiveness is typically confined to critical or well-understood components [4]. In complex engineering systems, interactions among subsystems and an incomplete knowledge of failure mechanisms often result in physics-based models that deviate significantly from real-world degradation dynamics [134].

**Table 7**

Checklist for Bridging the simulation-to-real gap in PHM.

Applications in PHM	Type of tasks	Fault detection; Fault diagnosis; RUL prediction.
	Addressed challenges	Simulation-to-real gap; Rare failure modes
Requirements	Prior knowledge	Physics-based models; First-principles simulations; Domain expertise
	Data	Synthetic data from simulations; limited real measurements.
	Type of bias	Observational bias
	Assumptions	Physics-based models capture relevant degradation mechanisms; Domain gap is manageable for the chosen approach.
Summary	Advantages	Reduced dependency on labeled real-world data; Coverage of rare failure scenarios
	Disadvantages	Simulator fidelity; Complexity in implementation and validation.
Resources	Repositories	ProgPy [143], Transfer Learning Library [144]
	Datasets	Physics-based simulator of the electrochemical battery [145] + Real Batteries [146]

Despite their potential, the adoption of physics-based models remains limited in practice [135], due to the complexity of their development and the expertise required for accurate parameterization. Recent advances have combined physics-based models with ML to alleviate the scarcity of labeled real-world data in PHM [136,137]. One widely used strategy is to generate synthetic training data from physics-based simulations. This methodology has been successfully applied to complex systems, such as aircraft engines [138] and roller bearings [139], to train classifiers where real failure data is practically non-existent. Furthermore, comprehensive reviews of Digital Twin technologies highlight the increasing role of these simulation-based strategies in modern modeling [140]. These synthetic datasets can serve multiple purposes, such as complementing sparse real-world measurements, creating examples of rare failure modes, and providing controlled variations of operating conditions. Typically, these physics-based simulations are used to model system dynamics and degradation processes, enabling researchers to generate comprehensive datasets that would be difficult or impossible to obtain from real-world operations. A major challenge in leveraging synthetic data for PHM is the Simulation-to-Real gap (Sim2Real), which is the discrepancy between idealized mathematical representations (simulations) and the complexities of real-world operating conditions. Even high-fidelity simulations cannot perfectly capture real-world variability, resulting in a distribution gap between synthetic data and experimental measurements of mechanical systems [141,142]. Table 7 provides a condensed checklist for applying Sim2Real strategies in PHM, outlining the primary applications, requirements, and trade-offs.

##### Applications of simulation-to-real approaches in PHM

To address the Sim2Real domain gap, researchers have developed strategies that leverage both physics-based simulations and real-world data [2,49]. These approaches aim to mitigate the distribution shifts inherent between simulated and real environments. The strategies differ in how they handle domain discrepancies and the requirements for labeled real-world data, with each approach offering distinct trade-offs between implementation complexity and data requirements. Fig. 6 summarizes three key strategies:

•**Approach 1: Data Augmentation** This strategy leverages physics-based models to enhance training datasets by enriching the information content available to machine learning algorithms. Data augmentation in the context of simulation-to-real transfer operates through two distinct mechanisms, each addressing different aspects of the data scarcity problem in PHM.

The first mechanism, *feature augmentation*, expands the feature space by fusing physics-derived quantities with observed sensor measurements. Rather than generating entirely synthetic samples, this approach computes additional features from physics-based models using real sensor data as inputs. For example, Arias Chao et al. [4] demonstrated this by fusing physics-derived performance parameters with sensor data to improve RUL prediction in turbofan engines. This strategy has proven effective in reducing the amount of real-world data required for training across various applications [147,148]. Because feature augmentation operates on real sensor data and only augments the feature representation, the primary source of uncertainty lies in the fidelity of the physics model rather than in the distribution shift between training and deployment domains.

The second mechanism, *sample augmentation*, addresses a different challenge by enlarging the training dataset through the injection of entirely simulated trajectories or degradation paths. This approach is particularly valuable in PHM contexts where real run-to-failure data are scarce or where specific fault modes are rarely observed in practice. By generating synthetic degradation sequences that span diverse operating conditions and failure modes, sample augmentation ensures broader coverage of the operational space. However, unlike feature augmentation, this strategy introduces a domain shift between simulated and real measurements, as the generated samples originate entirely from the physics-based simulator rather than from actual system observations. Consequently, the effectiveness of sample augmentation is more sensitive to the simulation fidelity and the magnitude of the sim2real gap. When the domain gap is substantial, sample augmentation may require explicit mitigation strategies such as domain randomization or hybrid data generation approaches that inject simulated fault signatures into healthy real-world signals [149].

•**Approach 2: Transfer Learning** This approach involves pre-training machine learning models using large-scale simulated data from first-principles simulations before fine-tuning them on a smaller set of labeled real-world measurements. This allows the model to learn fundamental physical behaviors from simulation while adapting to real-world noise and specificities through fine-tuning. This is particularly valuable when historical time-to-failure trajectories are scarce, as shown by Biggio et al. [150]. for Li-ion battery discharge prediction. Similar benefits have been observed in reliability analysis and wind turbine monitoring [151,152]. Similar to the data augmentation approach, the transfer learning approach improves generalization and reduces the need for large labeled datasets.

•**Approach 3: Domain Adaptation** This approach aims to align distributions between simulated and real-world data, enabling models to generalize across domains despite the absence of labeled data in the target domain [153,154]. Unlike transfer learning, which relies on labeled data for fine-tuning, domain adaptation learns domain-invariant features to bridge the distribution gap. Strategies such as those proposed by Liu et al. [153] and Wang et al. [154] learn domain-invariant features to bridge this gap. By improving the consistency between training and deployment distributions, these methods enhance model reliability in real-world scenarios where failure labels are unavailable.

#### *Advantages and limitations of simulation-to-real approaches in PHM*

The integration of physics-based data into data-driven methods offers distinct advantages for PHM, primarily by improving generalization to unseen operating conditions while reducing the need for a large number of labeled real-world data [155,156]. This benefit is critical given the cost and complexity of acquiring run-to-failure datasets for indus-

trial systems[157]. Studies have consistently shown that incorporating synthetic data can help model performance, as shown in [139] for bearing dynamics.

In practice, this beneficial form of observational bias is introduced deliberately by incorporating physically realistic synthetic data into the training process. A straightforward implementation uses high-fidelity numerical models to replicate fault behaviors or, when degradation models are available, to simulate progressive degradation, potentially covering rare or extreme operating conditions [19]. However, accurately modeling progressive degradation remains particularly challenging; consequently, most studies leverage physics-based simulations primarily to generate diverse “healthy” samples across varying operating conditions rather than to model gradual fault evolution. A second limitation is that models typically treat components in isolation, ignoring their inter-dependencies. To mitigate this, alternative strategies generate hybrid data by injecting simulated anomalies into healthy real-world signals [149], thereby preserving relevant operating conditions and reducing the domain gap between simulated and real domains.

Additionally, generating data from high-fidelity numerical models (e.g. using the finite element method) can be computationally demanding. Even a single simulation of a degradation process, such as fatigue crack growth, may require several hours to complete [158,159]. To overcome this bottleneck, researchers increasingly rely on computationally efficient surrogate models [160] (see part I of the review for more details [6]). These surrogate models serve as accurate proxies for the full-order simulations, enabling rapid analysis and prognostics without the exorbitant computation cost.

While data augmentation can significantly improve model performance, its effectiveness depends on both the fidelity of the physics-based simulations and their relevance to the specific classification or regression task. This highlights a key assumption in bridging the simulation-to-real gap: the extent to which the simulated phenomena accurately reflect real-world degradation processes in target application. However, systematic studies assessing the maximum tolerable domain gap for each strategy remain limited.

#### *Unexplored or emerging applications of simulation-to-real approaches in PHM*

A promising strategy for bridging the simulation-to-real gap in PHM is the use of large-scale foundation models, neural networks pre-trained on extensive, heterogeneous corpora that learn general representations and can be adapted to downstream tasks via fine-tuning, few-shot, or even zero-shot transfer [161,162]. Foundation models have shown remarkable success in different scientific domains, such as protein structure prediction [163], advanced weather forecasting [164,165], and their emerging applications in scientific machine learning [166]. The key advantages of foundation models stem from their two-phase approach. First, *pretraining*, where a single large-scale neural network learns to capture relevant patterns and structural relationships from extensive, diverse datasets. Second, *fine-tuning*, which enables the model to leverage these learned representations to excel at new tasks and adapt to different domains with limited labeled data [167].

Recent operator-transformer foundation models for scientific data (e.g. POSEIDON [168]) show that a network can be pre-trained on a limited set of PDEs and later zero-shot generalise to *unseen* equations, provided the underlying physics share structural similarities. Translating this idea to PHM, one could pre-train on large-scale, multi-condition bearing simulations and then adapt to real-world operating conditions by fine tuning on a small amount of real data. A comparable strategy has been recently applied in bearing dynamics, where a model trained on simplified lumped-parameter spring-mass simulations is adapted to high-fidelity multiphysics simulations of similar bearings, combining finite element analysis with elasto-hydrodynamic lubrication, using only a small number of cases. This illustrates how pretraining on simplified physics can be followed by targeted fine-tuning toward more

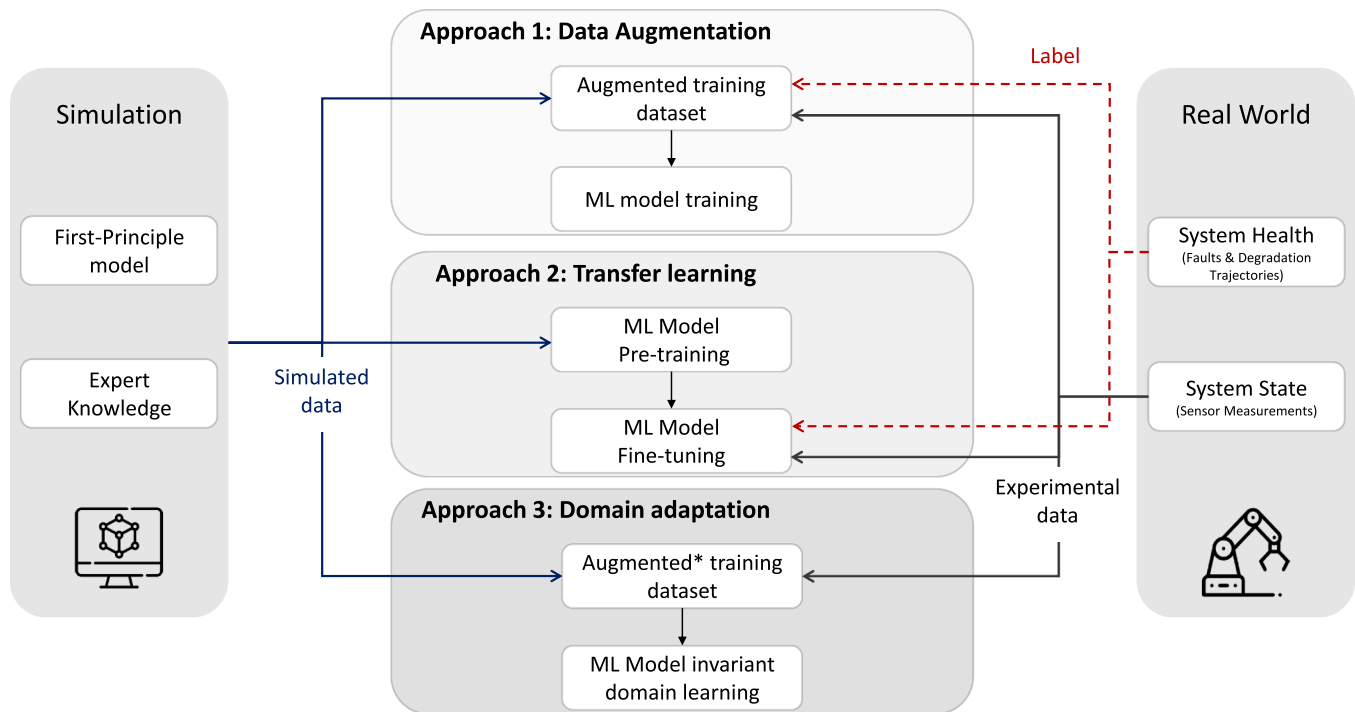


Fig. 6. Overview of approaches for bridging the simulation-to-real gap in PHM. The figure illustrates three primary strategies: (1) Data augmentation, which enriches training data by combining simulated and experimental measurements; (2) Transfer learning, where models are pre-trained on synthetic data and fine-tuned on real-world data; and (3) Domain adaptation, which learns domain-invariant features to align synthetic and real-world distributions. black arrows indicate simulated data, black arrows represent experimental data, and red dashed arrows denote labels.

realistic dynamics, consistent with the foundation-model paradigm discussed above [169].

Looking forward, a critical direction for methodology development is LLM-guided domain randomization. Inspired by recent advances in robotics, this approach uses foundation models to automatically optimize simulation parameters (e.g., friction or noise levels), ensuring that synthetic training data covers the real-world operating distribution without requiring manual tuning.

Beyond fine-tuning, these foundation models also offer potential for in-context learning [170–172] capabilities (see Section 5.1), where the model could analyze new degradation patterns without explicit retraining or fine-tuning them to real-world data. This enables rapid adaptation to novel failure modes or operating conditions simply by providing relevant examples within the input context.

### 3.3. Generative modeling

#### Introduction to generative modeling

Deep generative models (DGM) are neural networks designed to approximate complex, high-dimensional probability distributions using sample data [173]. These models estimate the likelihood of observed data and generate new samples that follow the learned distribution. With their high expressiveness, DGMs effectively capture complex data patterns, making them valuable for enhancing generalizability, reducing dependence on large labeled datasets, and generating rare samples to improve model robustness. Despite challenges such as computational inefficiency, the synthetic-to-real domain gap (Section 3.2), and the risk of bias propagation, DGMs remain a powerful tool for data augmentation.

Various deep generative models have been developed to learn data distributions. Among the most well-known are generative adversarial networks (GANs) [174] achieve this by introducing two competing neural networks: a generator, which learns to create realistic synthetic samples,

and a discriminator, which distinguishes between real and generated samples. In contrast, Auto-encoders (AEs) [175] take a different approach by learning to reconstruct original inputs while compressing the data into efficient latent representations, thereby capturing the essential features of the underlying distribution. A widely used variant, variational autoencoders (VAEs), introduce a probabilistic framework for encoding and decoding, allowing the model to learn a smooth and continuous latent space that facilitates both reconstruction and generation of new samples [176]. Flow-based generative models, such as those introduced in [177], take yet another approach by explicitly modeling the data distribution through a sequence of invertible transformations. More recently, diffusion models [178], inspired by non-equilibrium thermodynamics, have gained significant attention. These models define a Markov chain that incrementally adds noise to data during training, and then learns to reverse this process, reconstructing high-fidelity samples from pure noise.

**Physics-informed Data Generation** Despite their effectiveness in data generation, DGMs often require large datasets and may not always align with underlying physical laws. To improve their alignment with real-world phenomena, an emerging research direction focuses on integrating physics prior knowledge into generative models by embedding physical laws and invariance properties directly into the model architecture or training process.

One fundamental approach, as discussed in the Sim2Real section (Section 3.2), is to leverage physics-based simulations to generate training data for generative models. This strategy helps ensure that the generated datasets retain key physical properties with controlled variations while augmenting insufficient real-world data [137–139,186]. Another strategy is to interpret or encode data in terms of physically meaningful variables, thereby improving the model's adherence to physical principles. For instance, rather than directly predicting the wireless channel signals, Böck et al. [187] employed a VAE to estimate signal transmission parameters, which were subsequently used to reconstruct the channel matrix.

**Table 8**  
Checklist of generative modeling.

Applications in PHM	Type of tasks	Fault detection; Fault diagnosis; RUL prediction.
	Addressed challenges	Data scarcity.
Requirements	Prior knowledge	Governing PDEs (optional).
	Data	Enough data samples to represent the data distribution.
	Type of bias	Observational bias.
	Assumptions	The underlying data distribution can be learned by the designed model.
Summary	Advantages	Generalizability and adaptability; Reduced dependency on large dataset; Rare sample generation.
	Disadvantages	Computational inefficiency; Synthetic-real domain gap; Bias propagation.
Resources	Repositories	Deep Generative Modeling [179], GAN [180], TimeGAN [181]; Time-series Generative Adversarial Networks, MAD-GAN [182]; Multivariate Anomaly Detection with GANs, PyOD [183]; a library that includes VAEs and GAN-based anomaly detectors, TSDiff (Time Series Diffusion) [184]; A probabilistic diffusion model time series data imputation and forecasting
	Datasets	NASA C-MAPSS [51] and N-CMAPSS turbofan degradation datasets[52], PRONTO heterogeneous benchmark [185]

Beyond physics-informed data simulation and interpretation, researchers have developed methods to directly integrate physics into the training process of generative models. These approaches aim to enhance model reliability, interpretability, and generalization ability by embedding physics in different ways, ensuring consistency with the underlying laws of physics. One prominent research direction involves incorporating physics constraints into the loss function to enforce adherence to the governing equations. This approach, pioneered by PINNs (Section 2.1), uses physical laws as additional learning objectives during training. This framework has been extended to DGMs, for example Zhou et al. [188] proposed a physics-informed GAN for system reliability assessment, which fuses limited measurement data with underlying mathematical models that aggregates the system state probability in evolution of time. Their method demonstrated effectiveness on a simulated dual-processor computing system, highlighting its potential in scenarios where real-world failure data is scarce. In addition to physics-informed loss functions, researchers have explored integrating physical principles directly into the architectural design of generative models, such as graph neural networks (GNNs) for physical reasoning [189], Group equivariant Convolutional Neural Networks (G-CNNs) to model physical symmetry [39] and Hamiltonian Neural Networks (HNNs) [190]. With the advancement of diffusion models, Bastek et al. [191] introduced Physics-Informed Diffusion Models (PIDMs), which integrate physical constraints into the training of denoising diffusion models, enabling the generation of samples that comply with governing PDE equations. Along similar lines, Qiu et al. [192] applied PIDMs for fluid dynamics, proposing Pi-fusion – a method that integrates a physics-informed guidance sampling strategy into the diffusion process to capture quasi-periodical patterns in fluid motion, thereby demonstrating the potential of PIDMs for complex physical systems. Furthermore, Rissanen et al. [193] presented a generative modeling framework inspired by the physics of heat dissipation, leveraging the inverse process of heat diffusion to generate data that closely matches the true distribution.

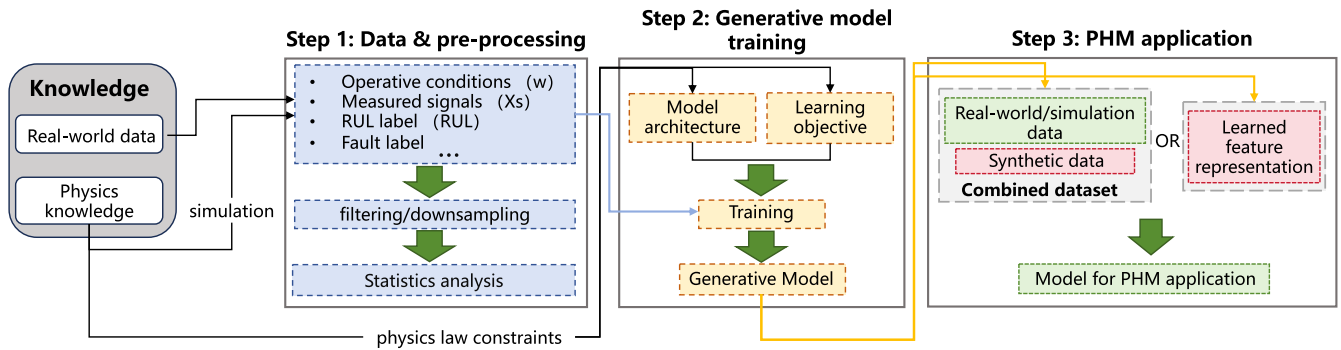
### Application of generative modeling in PHM

A major challenge in PHM applications is the scarcity of labeled training data in real-world industrial settings, as discussed in the previous subsection on simulation-to-real (Section 3.2) and in the section of scaling beyond single systems (Section 5). In response to this challenge, generative modeling has become an increasingly popular strategy for learning data distributions and generating realistic synthetic data in a flexible and efficient manner, as illustrated in the general framework shown in Fig. 7.

In scenarios characterized by data scarcity or imbalance, generative models are widely used to augment datasets and compensate for missing or underrepresented distributions. For instance, VAEs can learn the underlying data distribution and generate additional samples for minority classes, resulting in more diverse and balanced datasets. This approach has proven effective for fault diagnosis of bearing rolling [194–196] and multi-phase flow process [196] under imbalanced conditions. It has also been applied for RUL prediction of turbofan aircraft engines [197]. Beyond simply enriching data distribution, generative models are increasingly applied to synthesize rare faulty data, which is typically difficult to obtain in industrial systems. Hong et al. [198], for instance, combined VAEs with GANs to augment noise and vibration data, enhancing the representation of rare fault conditions and improving the accuracy of bearing fault diagnosis. Furthermore, generative models have been applied to address scenarios involving unseen fault conditions by simulating novel fault types, thereby improving the generalization capability of diagnostic models. For instance, Rombach et al. [199] utilized GANs to generate synthetic data representing previously unseen fault types under specific operating conditions, thereby enabling partial open-set domain adaptation. Sun et al. [200] deployed diffusion model to generate visual unseen anomalies for the unsupervised anomaly detection task. These approaches effectively address the challenge of limited real-world fault data by enriching the dataset with diverse fault scenarios.

Beyond data augmentation, generative models are increasingly applied to learn the underlying data distributions during the generative training, enabling the extraction of compact and informative representations. For instance, Ren et al. [201] introduced a multivariate invertible deep probabilistic framework to uncover latent patterns in gearbox fault diagnosis. By distilling essential low-dimensional features from high-dimensional sensor data, their method achieved more robust and accurate fault classification. In addition to feature learning, DGMs are leveraged to model the normal behavior of complex systems. This approach is particularly valuable for fault detection in scenarios where faulty data is scarce or unavailable. By learning the distribution of normal system behavior, these models can identify deviations as potential faults. For instance, VAE-based generative models have been used for time series reconstruction in wind turbine anomaly detection, successfully identifying deviations from expected signal patterns [202]. Hu et al. [203] proposed a diffusion-based weakly-supervised approach for deriving health indicators of rotating machines, by training a diffusion model to generate normal samples. Similarly, Zhu et al. [204] combined the temporal modeling capabilities of LSTM networks with the generative power of GANs, enabling the capture of complex temporal dependencies and dynamic behaviors. The fused model design significantly enhances fault detection by combining the generator's reconstruction error (the residual) and the discriminator's classification loss.

Physics-informed data generation has been increasingly applied in PHM to achieve more interpretable and reliable model performance by integrating known physical principles as constraints within generative models. For example, Gao et al. [141] introduced a GAN-based approach to generate synthetic training data that complements physics-based simulation outputs. By increasing data diversity, their method improves the generalization ability of models for bearing fault detection. Expanding on this idea, Nguyen et al. [186] proposed a fuzzy GAN that embeds



**Fig. 7.** General framework of generative modeling for PHM applications. The process consists of three main stages: (1) Data collection and pre-processing: collect real-world data, optionally augmented with physics-based simulated data, followed by data pre-processing; (2) Generative model training: train generative models to learn the underlying data distribution of input data, incorporating physical priors where applicable; and (3) PHM application: use the trained model to generate synthetic data, extract informative features to support downstream tasks such as fault diagnosis, anomaly detection, and remaining useful life prediction.

physical degradation models for RUL prediction to produce degradation trajectories that accurately mimic real-world system behavior, resulting in more robust RUL prediction for bearing systems and C-MAPSS dataset. To further integrate physical principles into the generation process, Fan et al. [205] first performed feature extraction to capture relevant physical characteristics, which were then used as reference vectors in a GAN-based model. This approach facilitated the generation of synthetic samples tailored to infrared images, showcasing how physics-informed feature extraction can improve diagnostic accuracy. Another widely adopted strategy in PHM involves incorporating physics-based constraints directly into the loss functions of learning models. Xu et al. [206], for instance, developed a physics-informed dynamic depth autoencoder that integrates battery state and capacity equations as penalty terms in its loss function. This design improves both the accuracy and interpretability of the lithium battery secondary variables, capacity, and state-of-health prediction results. Similarly, Xiong et al. [207] introduced the Controlled Physics-Informed GAN (CPI-GAN) for RUL prediction of turbofan engines under conditions of limited time-to-failure data. CPI-GAN generates synthetic trajectories of operating conditions that integrate degradation by incorporating five basic physics-based constraints as controllable terms within the loss function to enforce the generator. These constraints enforce consistency between the generator's outputs and the underlying physical laws, thereby enhancing both the performance and reliability of the RUL predictions. In terms of architecture design, PG-GAN [208] allows incorporating physical information within the architecture and training process of GAN. This enhanced GAN-based architecture uses a VAE to structure its generator's input, an ESO to provide its discriminator with real physical data, and a physical loss function to improve the generator's training. Sanchez et al. [209] augment sensor data using Physically Informed Echo State Networks (ESNs). Their approach integrates domain-specific physical knowledge into the learning process to generate surrogate time-amplitude signals from the Power Spectral Density of the data and is tested on the real-world jet fan dataset. A summary of the application contexts for generative modeling is provided in Table 8. Here, we outline the requirements and provides a summary of generative modeling. We review publicly available repositories implementing mainstream architectures, including Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), and Diffusion models. Furthermore, we highlight exemplar datasets widely utilized for generative tasks within the PHM domain. These datasets contain realistic sensor data collected under diverse operating conditions. Such heterogeneity is critical for evaluating the generalizability of generative models, enabling rigorous testing of their ability to synthesize robust data across varying domains and operating modes.

### Advantages and limitations of generative modeling in PHM

Generative models are playing an increasingly important role in PHM applications by addressing critical data limitations inherent to the field, including the scarcity of faulty samples, incomplete coverage of possible fault types, and the limited availability of complete time-to-failure trajectories under diverse operating conditions. A key advantage of these models is their ability to generate synthetic data to overcome these limitations. By sampling from low-density regions of the learned data distribution, they are able to generate more data for rare or underrepresented failure scenarios. This not only improves the representativeness of the dataset but also helps to balance heavily skewed datasets. Specifically, these models can generate new examples of rare failures to address the scarcity of faulty samples, and they can extrapolate from the learned distribution to generate unseen fault types, tackling the problem of incomplete fault coverage. For prognostics, these models can synthesize entire time-to-failure trajectories, providing the complete degradation data needed to train robust models for RUL prediction. Beyond data augmentation, generative models excel at learning robust representations of normal system behavior. Through training on data generation tasks, they learn to capture the intrinsic patterns of healthy operation across all conditions. This directly addresses the challenge of highly varying operating conditions, as the models can learn to distinguish between changes due to system operation and those indicating actual degradation. This capability greatly improves the generalization ability of the fault detection system, making it highly effective in detecting anomalies or deviations that may indicate potential faults.

However, applying generative models in PHM applications also comes with challenges. These models often require high computational resources and large amounts of high-quality training data, which may not always be available in real-world scenarios. Additionally, their complexity makes them difficult to interpret, which can be a drawback in safety-critical applications. Without proper regularization and sufficient data, these models also risk overfitting to the observational biases within the limited dataset, leading to biased learned distributions that fail to generalize to new conditions. In addition, the synthetic samples produced by these models may not always align with real-world physical laws, leading to potential inaccuracies. Despite these limitations, generative models remain an active area of research for improving predictive maintenance and fault detection.

### Unexplored or emerging applications of generative modeling in PHM

While generative models have demonstrated significant potential in PHM, many applications remain unexplored or are only beginning to

emerge. One promising direction is the integration of physics-informed constraints into generative models. By embedding prior physical knowledge, these models can generate data that adheres to the fundamental principles governing industrial systems. Another emerging area of development is the data-efficient generative modeling, which aims to improve the effectiveness of these models in real-world PHM applications in scenarios characterized by limited or imbalanced data and constrained computational resources.

#### *Promising research direction for PHM: Causally-informed synthetic data generation*

In addition to integrating physical laws, embedding causal knowledge into the generative process is an emerging and highly promising direction for generating reliable synthetic datasets. Just as physics-informed models introduce inductive biases that encourage learned representations to be consistent with physical principles, causal models bias generative models toward reproducing the true relationships between measured signals rather than confounding correlations. This causal inductive bias is typically integrated by embedding a system's causal graph directly into the generative architecture. Such graphs can be constructed using expert knowledge or obtained through causal discovery methods (see part I of the review [6] for more details). As a result, the synthetic data generated will reflect the underlying causal mechanisms instead of replicating spurious associations, leading to more trustworthy and interpretable datasets. While the discovery and use of such causal graphs as inductive bias were discussed in Part I [6], leveraging them specifically for *generative* purposes represents a promising next step for implementing observational bias in a principled manner.

The effectiveness of causally-informed synthetic data generation has already been demonstrated across several applications. For example, the tabular foundation model Tabular Prior-data Fitted Network (TabPFN) [210] was trained exclusively on synthetic data generated from causal models, demonstrating that models trained purely on causally-generated synthetic data can match or exceed the performance of models trained on real data for downstream tasks. More recently, first studies have begun to explore this direction in industrial settings. For example, [211] successfully employed causal models to generate synthetic samples for injection molding processes. By first discovering the causal graph from sensor data and then refining it with expert knowledge, they constructed a causal model capable of generating synthetic samples that closely matched real distributions.

For sensor-monitored dynamical systems where many variables are typically strongly correlated, causal generative methods offer a principled approach to generating data that is not biased by spurious correlations. Beyond this, causal methods make it possible to generate counterfactual examples and rare-event samples that correlation-based approaches typically fail to capture. Although their use in PHM remains limited, there is growing theoretical work on causally informed generative modeling for time series. One example is the normalizing flow network DoFlow [212], which generates counterfactual time series by defining the flow-based generative process over the system's causal graph. For instance, DoFlow could generate 'what-if' scenarios showing how turbine vibrations would affect power generator output in a power plant. This causal bias leads to strong generative performance on both synthetic time series data and two real-world case studies: time series of a hydropower system and cancer treatment outcomes.

Despite these opportunities, the widespread adoption of causal generative models faces practical hurdles. First, causal architectures introduce additional complexity, as they must incorporate the inductive bias implied by the causal graph. Second, while numerous methods exist to infer causal graphs from data, integrating causal discovery into the modeling workflow adds yet another layer of complexity. Moreover, causal discovery alone is often insufficient: The resulting graph frequently requires refinement through expert knowledge. A promising development in overcoming these limitations is the creation of bench-

marks and datasets with known causal graphs, such as [213]. Momentum in this direction is further supported by the growing adoption of causal chamber experiments [214], an emerging experimental platform in which the true causal graph is known, enabling datasets to be paired with ground-truth causal structures.

In summary, progress in causal datasets, together with advances in causally-informed generative models for time series, is laying the foundation for incorporating causal observational bias into PHM. Challenges remain, particularly the complexity of causal methods and the limited availability of reliable causal graphs, but ongoing developments suggest that causal generative models may soon enable more accurate, trustworthy, and interpretable synthetic datasets for PHM.

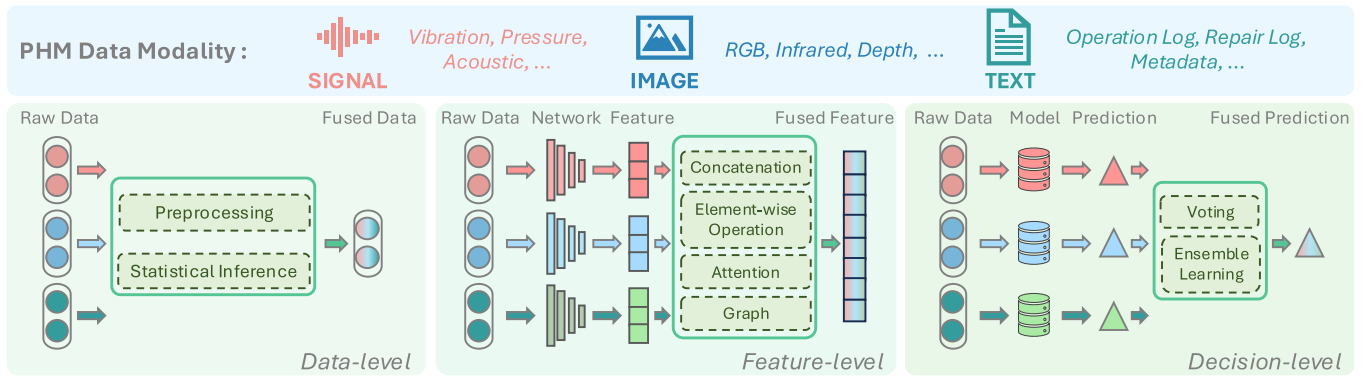
### 3.4. Data fusion

#### *Introduction to data fusion*

Data fusion refers to the integration of information from diverse sources and modalities into a unified representation [215,216]. By leveraging cross-domain data, it can capture different partial information about the same system condition and compensate for the limitations of individual sources, thereby improving the overall accuracy, robustness, and reliability of downstream algorithms [217,218]. The ability to integrate complementary information, reduce uncertainty, and improve decision-making has driven the widespread adoption of data fusion techniques across various domains. Notable applications include autonomous vehicles [219,220], computer vision [221,222], and wearable health monitoring systems [223]. In PHM, the increasing complexity of industrial systems makes it challenging for uni-modal data to accurately describe overall degradation processes. As a result, data fusion techniques enable the integration of additional heterogeneous data streams (e.g., temperature, pressure, vibration) to support fault detection, diagnostics and prognostics in systems like aircraft engines [224,225], gearboxes [226], and electric vehicles [227].

From a PIML perspective, multimodal data fusion inherently introduces observational bias, not merely as noise or artifacts, but as a set of constraints imposed by the selection, representation, and sampling of input data [228]. In practice, the choice of sensors, included modalities, spatial configuration, and underlying measurement principles collectively determine which aspects of the physical system are emphasized or omitted in the fused representation [229]. Rather than being a direct reflection of the physical system, the integrated data is shaped by the sensitivities and limitations of each modality, creating a composite representation defined by both the complementarity and incompleteness of the data sources. In PIML, these observational biases influence model learning by shaping which physical laws or latent variables can be inferred from the available observations [218,230]. While such biases can ensure that models remain grounded in sensor-accessible, physically relevant information, they also risk obscuring or distorting less observable phenomena. Explicitly acknowledging and managing this bias is therefore essential for PIML as it enables models to leverage the strengths of data fusion for robust and interpretable inference, while mitigating the risk of spurious correlations or information gaps that may compromise physical reasoning and prediction.

Data fusion methodologies are commonly categorized by the characteristics of data sources involved: heterogeneous and homogeneous fusion. Heterogeneous data fusion combines information from diverse modalities or sensor types. In PHM, commonly used data include vibration, acoustic, temperature, pressure, speed, electric current, and voltage signals, with visual and text data being increasingly used in recent years [218,230]. However, significant challenges remain due to inherent variations in data quality, structure, and representation formats [228]. In contrast, homogeneous data fusion integrates data from multiple sources of the same modality. For instance, multi-camera systems positioned at *different locations* capture complementary visual perspectives of the same scene, supporting applications such as building inspection [231,232] and rail defect detection [233,234].



**Fig. 8.** Overview of data fusion strategies in PHM applications. This includes three primary fusion stages: (1) **Data-level fusion** combines raw sensor data through preprocessing or statistical inference to generate integrated data representations; (2) **Feature-level fusion** extracts modality-specific features and integrates them using advanced techniques; and (3) **Decision-level fusion** merges predictions from individual models trained on different modalities using strategies like voting or ensemble learning.

Recent advancements in multimodal machine learning have broadened the scope of data fusion in two key directions: (1) the development of novel neural representations by integrating common modalities, such as combining infrared images with vibration signals for gearbox diagnosis [235]; and (2) the incorporation of previously underutilized modalities such as natural language and semantic metadata, for example, integrating technical documents and maintenance logs into diagnostic models for hydrogenerators [236,237]. To systematically categorize recent data fusion methodologies, we adopt a common classification framework based on three fusion stages [216,238]: **data-level fusion**, **feature-level fusion**, and **decision-level fusion**, as illustrated in Fig. 8. Each level presents distinct trade-offs in terms of complexity, interpretability, and performance.

**Data-level fusion** focuses on the direct combination of raw data streams to estimate system states or conditions [216,239,240]. Classic approaches are typically classified into two complementary groups: preprocessing techniques and statistical inference models. Preprocessing techniques, such as image processing and signal filtering [241–243], operate directly on raw data to achieve spatial, temporal, or statistical alignment of multimodal inputs. These methods preserve the *original data format* while enhancing accuracy and minimizing noise, thereby ensuring that downstream models operate on well-calibrated, high-quality input data.

In contrast, statistical inference methods, including the Kalman Filter, Extended Kalman Filter, Unscented Kalman Filter, and Particle Filter, recursively integrate raw sensor observations to estimate the state of dynamic systems (*prediction stage*) while compensating for measurement and process noise (*correction stage*) [244]. Further details on the use of Kalman filtering for state transitions and sequential modeling are provided in the State-Space Models section of Part I.

**Feature-level fusion** extracts high-dimensional representations from preprocessed data and performs cross-modal integration to construct more discriminative feature representations [245]. In the context of feature-level fusion, traditional feature extraction techniques used in ML-based PHM applications, such as Fourier and wavelet transforms [242], are typically designed for a single modality. As a result, they are generally inadequate for fusing multimodal inputs that differ in sampling rates, physical units, or data formats (e.g., time-series signals, images, or tabular measurements) [246,247]. In contrast, recent approaches leverage deep learning architectures that automatically *learn* relevant features from data. These architectures are typically applied independently for each modality to effectively capture modality-specific information [240,248]. To fuse these modality-specific representations, various fusion strategies have been proposed. Basic methods, like concatenation or element-wise operations on learned features, are widely

used [249,250] but often fail to capture inter-modality correlations, potentially retaining redundant information. Recent studies incorporate techniques such as Canonical Correlation Analysis or attention mechanisms to emphasize complementary information while suppressing irrelevant or overlapping features across modalities [240,251]. Moreover, graph-based models, including Graph Neural Networks, Graph Convolutional Networks, and Graph Attention Networks, have emerged, enabling the modeling of complex interdependencies and contextual relationships within multimodal data [240,252].

**Decision-level fusion** combines predictions from multiple independent single-modality models to reach a final decision [245,253]. Each modality-specific model generates predictions based on its respective measurements and algorithms, thereby offering a high degree of flexibility in accommodating heterogeneous data characteristics and maintaining robustness even in the presence of sensor noise or failure [254]. These methods are typically categorized into statistical approaches and learning-based techniques. Statistical methods use probabilistic or rule-based frameworks, such as weighted voting [255], Dempster-Shafer theory [256], and Bayesian decision theory [257], to reconcile uncertainty and inconsistent decisions from different modality-specific models. In contrast, learning-based methods often adopt ensemble learning strategies, where independent sub-models are trained on each modality to estimate predictive confidence or contribution weights. The final decision is then obtained by aggregating these outputs, typically through confidence-weighted averaging or voting, to improve overall reliability and accuracy [258,259].

#### Application of data fusion in PHM

Benefiting from advances in machine learning and the increasing availability of multimodal data, data fusion has become a critical enabler in PHM, driving advancements in fault detection, diagnosis and prognostic through the integration of multi-source and multi-modal information [218,247]. A concise overview of representative application contexts, addressed challenges, and typical requirements for data fusion in PHM is provided in Table 9. Here, we also include Bayesian filtering implementations and pretrained multimodal architectures that are particularly relevant to data-level and feature-level PHM applications. Moreover, we summarize a set of widely used public datasets that underpin much of the existing literature on data fusion in PHM.

At the data level, **preprocessing techniques** serve as a preliminary step to ensure high-quality, well-calibrated input data for subsequent feature- or decision-level fusion. For example, frequency-domain transformations and cosine similarity are commonly applied to align and merge homogeneous vibration signals collected from multiple sensor

**Table 9**  
Checklist of data fusion.

Applications in PHM	Type of tasks	RUL prediction; Fault detection; Fault diagnosis.
	Addressed challenges	Uni-modal bias; Incomplete health representation; Degradation modeling.
Requirements	Prior knowledge	Sensor models; Modality-specific representation.
	Data	Heterogeneous or homogeneous data.
	Type of bias	Observational bias.
	Assumptions	Different modalities are interconnected with shared complementary information.
Summary	Advantages	Cross-modal data compensate for individual weaknesses and improve overall system accuracy and reliability.
	Disadvantages	Inherent differences in the statistical properties of each modality and the complex nonlinear relationships among their raw representations.
Resources	Repositories	FilterPy <sup>a</sup> , Hugging Face Transformers <sup>b</sup>
	Datasets	NASA C-MAPSS / N-CMAPSS, NASA Li-ion Battery Aging, PHM Society Data Challenges, etc.

<sup>a</sup> <https://github.com/rllabbe/filterpy>

<sup>b</sup> <https://huggingface.co/docs/transformers/index>

positions. The fused signals are then fed into feature-level architectures for fault detection, either with CNNs combined with acoustic modalities in compressor blades [260], or independently into neural networks for rotating machinery [261]. Beyond fault detection, a common data-level application in PHM is the construction of **health indices (HIs)**. HIs are typically generated by fusing multiple sensor measurements to support degradation modeling and prognostics. Various approaches have been proposed for HI construction, ranging from quadratic programming-based fusion [262,263], genetic optimization [264], kernel-based fusion [265], and more recently, deep neural networks that predict HIs by capturing nonlinear relationships among multiple sensor signals [266].

In contrast to preprocessing techniques, **statistical data-level fusion** approaches, particularly the Kalman Filter and its variants, have been widely adopted in PHM. By leveraging probabilistic state estimation and recursive updates, KF-based methods combine established system models with measurements from multiple modalities to enable prognostics [267,268]. However, traditional KF applications are often limited by their reliance on linear system assumptions and the requirement for precise a priori models, such as system dynamics and noise characteristics. Recent advancements have focused on improving the adaptability and robustness of KF-based methods through integration with data-driven techniques for RUL prediction [269,270]. While the KF assumes linear system dynamics and Gaussian noise [271], the Particle Filter (PF) employs sequential Monte Carlo sampling to handle highly nonlinear and non-Gaussian systems [272]. For example, PF has been combined with SVM trained on fused HIs, where entropy-based fusion was employed to fuse the prognostic outputs from PF over multiple time steps [269]. Similarly, Lee et al. [270] utilized an EKF to fuse vibration and pressure signals, which were subsequently used to train a BiLSTM network for RUL prediction. These integrations leverage interpretability of model-based filtering and the adaptability of data-driven learning to capture nonlinear degradation patterns and complex system dynamics under time-varying conditions [273].

At the **feature-level**, most PHM applications have focused on **signal-based** modalities, reflecting the prevalence of vibration, acoustic, and other time-series measurements in condition monitoring. In this setting,

features are first extracted from each modality, either through hand-crafted signal processing or through learnable transformations, before integration. Common architectures include Fully Connected Neural Networks for learning nonlinear feature representations [274,275], CNNs for capturing spatial patterns [276,277], and RNNs, especially LSTM variants, for modeling temporal dependencies [278,279]. Recent advances have extended feature-level fusion to increasingly heterogeneous modalities, broadening the scope of PHM applications through the integration of **image and text** data. For example, in the visual domain, Cordoni et al. [280] concatenated pretrained VGG features extracted from thermal images with structured data (e.g., internal temperature, power consumption) through an autoencoder, forming a unified feature representation for refrigerator condition monitoring. To address feature distribution misalignment in cross-modal fusion, Zhang et al. [281] combined separately extracted features from thermal images and vibration signals through fully connected layers, jointly optimizing classification, maximum mean discrepancy, and cross-modal consistency losses for gearbox diagnostics. Similarly, the DFNet architecture incorporates a feature interaction module to enable cross-domain feature learning and a global fusion module that adaptively integrates shallow and deep representations from vibration, infrared, and spectral data for rotating machinery diagnosis [282]. For electrical power equipment inspection, capsule networks have been employed to fuse RGB, infrared, and ultraviolet imagery, leveraging vector-based features to preserve spatial relationships between low- and high-level features and jointly capture leakage, temperature, and other physical information [283].

**Text-based feature fusion** has gained increasing attention by integrating features from unstructured textual sources, such as maintenance logs, condition reports, and technical documents, with those from other sensing modalities, including time-series data and images. For example, Wang et al. [284] formulated fault diagnosis as a *visual question answering* task, employing large language models to fuse textual expert knowledge with envelope spectrum images constructed from vibration signals through the Hilbert transform and FFT. In this framework, visual features from envelope spectrum images are projected into the LLM embedding space and concatenated with text embeddings, enabling multimodal reasoning that improves bearing fault diagnosis accuracy, particularly in limited data scenarios [284]. Beyond this, domain-adapted LLMs have been coupled with condition monitoring data to extract contextual insights from maintenance logs in hydropower generators [237]. In this approach, the models are fine-tuned with domain-specific knowledge and fused narrative fault descriptions with sensor-derived operational parameters to enhance situational awareness. Similarly, Mandelli et al. [285] integrated natural language processing with model-based system engineering to fuse unstructured nuclear plant reports with structured system models, thereby linking causal and health-related information directly to component functions and dependencies.

GNNs have also emerged as a powerful framework for **feature-level fusion** in PHM. A representative class of approaches is spatial-temporal fusion architectures, which integrate graph-based physical topology (e.g., sensor/module relationships) with temporal sensor data. In such frameworks, GNNs model the spatial structure, while LSTM- or attention-based encoders capture temporal dependencies. This has been demonstrated, for instance, in turbofan engine RUL prediction using hierarchical sensor-module graphs with adaptive cross-graph fusion [286], as well as in hierarchical attention graph convolution with self-attention pooling for selective sensor integration [287]. Building on this, heterogeneous graph fusion methods explicitly represent complex system interactions by encoding diverse node and edge attributes. Examples include multiplex aggregation across spatial and temporal meta-paths to fuse heterogeneous bearing vibration patterns [288], as well as heterogeneous graph attention mechanisms that integrate both intra-domain and inter-domain dependencies (e.g., hydraulic-electrical subsystem couplings) with distinct temporal dynamics in hydropower plants [252,289]. In addition, hierarchical feature fusion techniques

capture multi-scale sensor correlations through multi-level attention mechanisms, such as node-, subgraph-, and cyclical-level modeling in aircraft engines [290], or adaptive weighting strategies applied to micrograph sequences [291]. Beyond this, several methods augment GNN-based fusion with **physics-informed or statistical regularization**, such as embedding autoregressive moving-average regression within GCNs or incorporating physics-based loss terms to prioritize timely fault prediction [292]. Other approaches compute performance deterioration indices by fusing residuals between predicted and measured heterogeneous signals in hydropower turbines [293]. These enhancements improve both the robustness and interpretability of PHM models. A more detailed introduction to GNN architectures and related PHM applications is provided in the Graph Neural Networks section of Part I.

#### *Advantages and limitations of data fusion approaches*

Data fusion techniques significantly enhance PHM systems by effectively integrating multi-source and multimodal information, thereby improving the accuracy, robustness, and adaptability of detection, diagnostic, and prognostic models. By combining diverse data sources, fusion methods provide a more comprehensive and complementary representation of system health and degradation process, capturing failure patterns that may be overlooked by uni-modal approaches. From a Physics-Informed Machine Learning perspective, incorporating complementary information from different modalities reduces observational bias and aligns learned representations more closely with underlying physical processes. This enriched, physically consistent view leads to more meaningful insights and improves the generalization of health monitoring models across varying operational scenarios.

A key advantage of data fusion lies in increased robustness and fault tolerance, achieved through redundancy and complementary information. This mitigates the impact of noisy, incomplete, or missing data, which ensures greater reliability and continuity under potentially adverse conditions. Moreover, data fusion enables the modeling of complex spatial, temporal, and structural dependencies often overlooked by single-source methods through advanced techniques like Transformers and GNNs, which are capable of capturing interactions across modalities and components. An additional benefit is improved scalability and adaptability through hierarchical and hybrid fusion architectures, which support multi-level integration and dynamic reconfiguration in response to changing system states. Finally, by incorporating heterogeneous modalities such as text and image data, data fusion broadens the informational scope of PHM systems. This allows for context-aware diagnostics and more interpretable predictions by combining quantitative sensor measurements with unstructured expert knowledge or visual evidence.

Despite its advantages, several challenges still limit the practical deployment and effectiveness of data fusion in PHM systems. From a data perspective, heterogeneity remains a significant obstacle. Sensor data from different modalities often vary in sampling rates, formats, temporal alignment, and noise characteristics. Integrating such heterogeneous streams requires extensive preprocessing and complex architectures for effective fusion. From a methodological perspective, a critical limitation lies in the limited interpretability of many fusion models, particularly those based on deep learning. These models often operate as black boxes, making it difficult to validate system behavior or capture cross-modal correlations effectively. When interactions between modalities are ignored or insufficiently modeled, the resulting fused representation may not leverage the richness of multi-source data.

#### *Unexplored or emerging applications of data fusion in PHM*

Modern data fusion in PHM is no longer limited to traditional physical sensors such as accelerometers, strain gauges, or temperature probes.

Increasingly, it incorporates **non-sensor data sources**, including maintenance records, inspection reports, fault annotations, and other semantic or contextual metadata, that provide complementary operational and historical insights.

Although the application of LLMs in PHM is still at an early-stage, aforementioned studies demonstrate their potential for **multimodal fusion**, where unstructured text is combined with structured numerical or visual sensor data. In these settings, the **text modality** typically encodes expert knowledge, historical fault cases, or procedural descriptions, whereas the **sensor modality** may consist of time-series measurements or derived representations such as envelope spectrum images. By mapping these modalities into a shared embedding space, LLMs can perform context-aware reasoning that links descriptive knowledge to observed machine states.

Further research aims to integrate 3D reconstruction (e.g., photogrammetry or neural radiance fields) and advanced computer vision (e.g., infrared anomaly detection) to construct digital twins or inspect visual degradation with high precision. These hybrid approaches bridge the gap between physical sensor data and contextual, human-generated knowledge, advancing PHM through enhanced interpretability and hybrid physical-contextual insights. Importantly, these advancements also highlight the need for models that can generalize across diverse operating conditions and data domains, a challenge that is increasingly central to the next stage of PHM research.

## 4. Reinforcement learning

### *Introduction to reinforcement learning*

Reinforcement Learning (RL) is a branch of machine learning in which an agent learns to make decisions by interacting with its environment [305]. The agent takes actions and receives feedback in the form of rewards or penalties, using this information to improve its strategy over time. The goal is to maximize cumulative reward by discovering optimal policies through trial and error. Unlike supervised learning, which depends on labeled examples, RL agents learn autonomously from feedback signals received as a result of their actions. This approach is especially well-suited for sequential decision-making problems where the optimal strategy cannot be predefined [306]. Recent advances in deep reinforcement learning have demonstrated impressive capabilities across diverse domains [307], including discovering a novel matrix multiplication method [308], mastering complex strategy games [309,310], controlling robotic systems [311], and advancing scientific discovery [312,313]. In the context of PHM, this ability to learn adaptive policies for complex, sequential decisions is particularly critical. While traditional PHM techniques focus on passive prediction, effective asset management requires translating these predictions into active decisions that directly impact the physical system. RL **closes this loop** by enabling agents to interact directly with their environment, take actions, and continuously refine maintenance and operational strategies in response to real-world feedback [314,315].

### *Applications of reinforcement learning in PHM*

PHM is evolving from a focus on predicting system conditions and remaining useful life to a more proactive paradigm: prescriptive operation and maintenance [316,317]. This emerging direction not only anticipates failures but also seeks to optimize system operations in real-time data, taking actions that extend asset life, reduce costs, and ensure safety. Achieving these goals requires approaches capable of making sequential, data-driven decisions in highly complex, dynamic environments. RL is particularly well-suited to these challenges due to its ability to learn from experience, adapt to evolving conditions, and optimize policies for long-term objectives. While research into prescriptive operation is still emerging, RL has already demonstrated promise in decision support for maintenance scheduling and optimization. In

**Table 10**  
Checklist of reinforcement learning.

Applications in PHM	Type of tasks	Prognostics; Fault detection; Prescriptive Operations
	Addressed challenges	Complex decision making tasks with long-term objectives and scarce supervision
Requirements	Prior knowledge	Environment interaction (potentially offline/simulated).
	Data	Depending on the task, many interactions with the environment are needed to obtain a robust internal representation of the task and learn a policy that accumulates satisfactory reward levels
	Type of bias	Learning, Inductive and Observational bias.
	Assumptions	The environment's dynamics can be learned by the agent through observations and interactions; Sparsity of the reward signal can be mitigated
Summary	Advantages	Multi-objective optimization; Sequential decision-making; Integration with simulation environments and digital twins for safe policy development
	Disadvantages	Data inefficiency; Hard to train; Hard to define a good reward signal; Scalability
Resources	Simulation environments	<b>RLfleet</b> : Fleet-level aircraft maintenance environment for deep RL [294]; <b>IMP-MARL</b> : Multi-agent RL benchmark for infrastructure maintenance [295]; <b>InfraPlanner</b> : Bridge network planning simulator for hierarchical RL [296]; <b>Industrial Benchmark</b> : Continuous-control benchmark inspired by industrial process operation [297]; <b>Grid2Op</b> : Power-grid operation and contingency management environment [298]; <b>CausalChamber</b> : Physical testbed for RL and causal reasoning [299].
	Repositories & toolkits	<b>IMPRL</b> : Modular library for inspection and repair RL scheduling [300]; <b>NASA PCoE / ProgPy</b> : Official prognostics models and tools for PHM-focused simulation and RL integration [301]; <b>Stable-Baselines3</b> : Reference RL algorithm implementations for rapid prototyping [302]; <b>OpenAI Gym / Gymnasium</b> : Standardized RL environment API and interface suite [303,304].

these applications, RL enables data-driven strategies that improve maintenance timing, resource allocation, and system reliability [318,319]. Recent studies have demonstrated RL's potential in domains such as energy systems [320], and manufacturing [321], underscoring its growing relevance for PHM. A particularly promising development is the rise of Physics-Informed Reinforcement Learning (PIRL), which provides systematic approaches for integrating physical knowledge into RL [322]. A summary of the application contexts for RL is provided in Table 10.

Building upon the PIRL framework, there are several concrete approaches for incorporating physical knowledge into RL systems. Physical priors can be integrated into RL using **safety filters** by refining RL agent actions to ensure operational constraints are never violated. In PHM, such filters can prevent RL agents from selecting actions that would cause excessive equipment wear, violate health limits, or trigger unsafe states during both exploration and deployment [320]. More advanced approaches leverage Control Barrier Functions (CBFs) to provide formal safety guarantees, ensuring that the learnable policy remains within the safe operational envelope defined by the physical limits of the system [323,324]. **Physics-Informed Neural Networks** [325] embed domain-specific knowledge, such as degradation models or conservation laws, directly into the policy or value function architecture. The Phyllis framework for data center cooling control [326], which incorporates thermodynamic equations into the RL pipeline, achieves faster and safer adaptation to changing operational conditions and improved energy efficiency compared to conventional methods.

Incorporating **task-related inductive bias** by embedding physical models or system structures within the RL environment further narrows the hypothesis space and guides agents toward physically plausible behavior. Model-based RL approaches are especially applicable to PHM, as they can integrate detailed system dynamics and degradation processes into the agent's world model [327]. This is illustrated by the reinforcement twinning framework, in which digital twins are co-trained with RL agents, and by the work of Bellani et al. [328], who demonstrated that embedding a physical degradation model and RUL estimation into an RL environment enables optimal operation and maintenance policies that outperform traditional rule-based strategies by adapting proactively to real-time degradation feedback [4,329]. Furthermore, **Graph-structured representations** offer another form of inductive bias, particularly for systems with complex inter-dependencies. For example, [330] modeled aircraft maintenance stand scheduling as a heterogeneous graph of maintenance stands and items, using graph

neural networks and deep RL to efficiently capture system relationships. This approach led to superior scheduling performance compared to traditional heuristics, illustrating the advantages of graph-based inductive bias in complex PHM decision-making [331].

In summary, the integration of physical knowledge and inductive biases into RL architectures marks a major step toward safer, more interpretable, and more effective PHM solutions, an overview can be seen in Fig. 9. These advances are paving the way for RL-driven maintenance and operational strategies that are robust to real-world complexities and capable of delivering both predictive and prescriptive value.

#### Advantages and limitations of reinforcement learning

RL offers several distinct advantages over traditional rule-based or supervised learning methods. System degradation in PHM is inherently complex and uncertain, making learning-based methods especially valuable for enabling proactive, adaptive operation. RL is inherently designed for sequential decision-making, enabling agents to optimize maintenance and operational actions over time and adapt strategies dynamically as system conditions evolve. Additionally, RL also supports multi-objective optimization, balancing competing goals – such as minimizing downtime, reducing maintenance costs, and maximizing safety or reliability – through carefully designed reward functions [332,333].

Moreover, RL can be combined with model-based approaches and digital twins, leveraging simulation environments to train agents safely before deployment [334]. This reduces real-world risk and enables robust sim-to-real transfer. Incorporating physics-informed representations and domain knowledge further enhances interpretability, data efficiency, and generalization, making RL highly effective for proactive PHM solutions.

However, several fundamental limitations currently constrain the widespread deployment of RL in PHM applications. One of the most significant challenges is sample efficiency: RL agents require extensive environment interaction, often impractical or unsafe in industrial settings, necessitating high-fidelity simulators or digital twins for offline training. However, simulators must accurately capture complex system behaviors and degradation mechanisms, and even then, there is a persistent risk of a sim-to-real gap, where policies that perform well in simulation may not generalize to actual operations (as discussed in Section 3.2). In this context, PIML approaches such as the surrogate models, PINNs, and GNN-based representations introduced in Part I-can be used

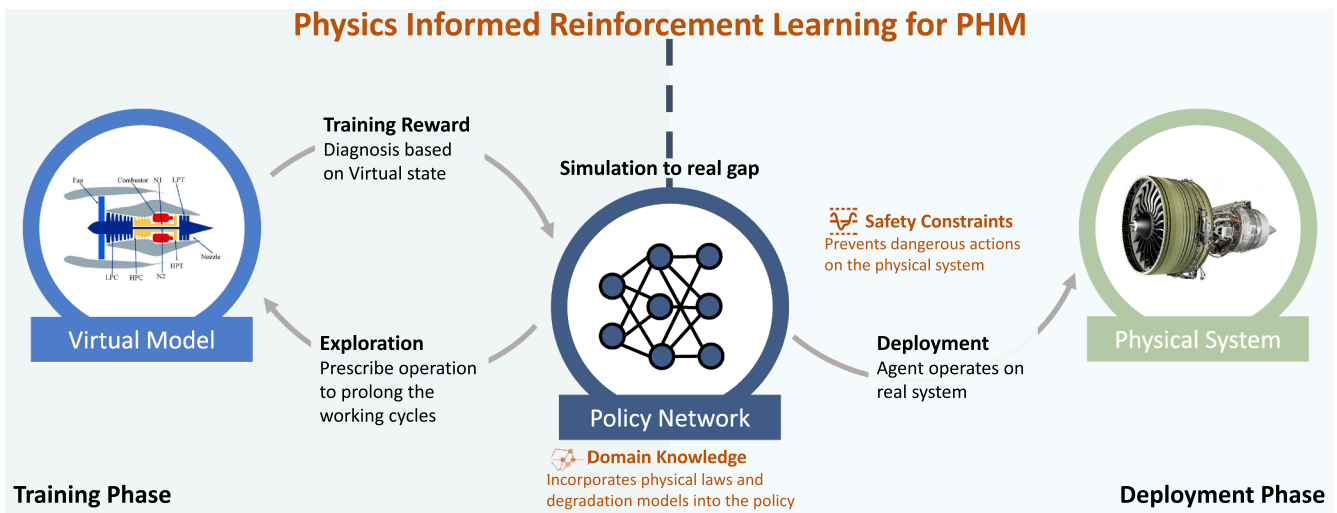


Fig. 9. The framework shows safe policy training in a virtual model (left) followed by deployment on physical systems (right), with physics-informed RL including safety constraints and domain knowledge integration. The simulation-to-real gap represents a key transfer learning challenge in real-world deployment.

to improve the fidelity of model-based RL simulation environments and reduce sim-to-real discrepancies by embedding physical structure and constraints [6]. Also, Offline Reinforcement Learning (Offline RL) avoids these safety risks by learning policies entirely from static, historical datasets. To mitigate distributional shift, where the agent overestimates the value of actions rarely represented in the dataset, algorithms such as Conservative Q-Learning (CQL) regularize the value function to be pessimistic about out-of-distribution actions [335]. Although CQL was not originally developed for maintenance optimization, its conservative design principles can be adapted to learn safe maintenance policies from logged operational data, complementing recent efforts in offline maintenance decision-making [336]. Beyond value-based offline RL methods, recent work has reformulated offline reinforcement learning as a sequence modeling problem using Decision Transformers [337], suggesting an alternative paradigm for learning maintenance policies directly from historical trajectories without explicit value estimation.

Another major challenge lies in the design of reward functions. Translating complex maintenance objectives—such as balancing cost, reliability, downtime, and risk—into a scalar reward signal is inherently difficult. Poorly specified rewards may incentivize unintended or even unsafe behaviors. To address this bottleneck, Inverse Reinforcement Learning (IRL) has emerged as a promising research direction. IRL inverts the problem by inferring the underlying objective function directly from expert demonstrations [338]. In PHM contexts, where expert operators often rely on implicit operational knowledge that is difficult to codify, recent IRL approaches can extract latent risk preferences from historical logs. This effectively allows for imitating expert decision-making policies by capturing implicit operational preferences, and may improve robustness to modeling and reward misspecification compared to manually designed reward functions in stochastic environments [339].

Beyond reward formulation, a further challenge is determining the optimal integration of physical prior knowledge. Imposing too many constraints can limit the agent's adaptability, whereas insufficient structure may lead to unrealistic behaviors. Finally, RL methods in PHM must contend with sparse rewards, partial observability, and non-stationary system behavior. For instance, the effects of maintenance actions often manifest only after long operational periods, and sensors may not fully capture degradation states. These factors complicate the learning process, highlighting the critical need for high-fidelity simulation tools and safety-aware exploration strategies.

#### Unexplored or emerging applications of reinforcement learning in PHM

Combining Physics-Informed Machine Learning with model-based RL [340] offers a promising frontier for PHM. By embedding physical laws and constraints directly into simulation environments, PIML enables the construction of more accurate and generalizable system models, a critical requirement for effective model-based RL. This synergy has already proven valuable in robotic systems, where PIML dynamics models allow RL agents to learn more effective policies and generalize across a broader range of operational scenarios [341]. Extending this approach to PHM could yield more robust, efficient, and trustworthy maintenance and operational strategies that adapt dynamically to the complexities of real-world assets.

Furthermore, differentiable physics simulators show great potential for sample-efficient, gradient-based optimization of maintenance and operational policies, as they expose analytic gradients of simulation outcomes with respect to control actions [342]. Gradient information let agents optimise long-term maintenance or operational objectives with orders-of-magnitude fewer samples than classic model-free methods, even when the underlying dynamics are highly nonlinear, as enabled by differentiable, physics-informed surrogate models reviewed in Part I [6].

Equally important is the choice of state representation. Augmenting or replacing conventional state representations with physical parameters, health indicators, and degradation metrics can make RL agents more sample-efficient and interpretable. Moreover, incorporating Multimodal RL allows agents to reason using heterogeneous data sources—combining time-series sensor readings with visual inputs or textual maintenance logs (as discussed in Section 4)—to build a more comprehensive state representation. Techniques addressing modality alignment and dynamic importance weighting [343] are particularly crucial here, ensuring that the agent learns to prioritize the most relevant sensor inputs for decision-making. In PHM, these physics-informed features, often derived from domain expertise or monitoring data, enable agents to focus on the most relevant aspects of system health and degradation. This approach not only accelerates learning and improves generalization but also enhances the transparency and trustworthiness of decision-making in practical maintenance scenarios.

Addressing the challenge of sparse rewards and complex, long-horizon objectives, methods such as Automatic Curriculum Learning (ACL) [344] and Hierarchical Reinforcement Learning (HRL) [345] provide structured mechanisms for training agents on progressively more challenging health-aware control problems. These approaches can

facilitate learning in PHM-related settings where agents must reason over degradation dynamics or multi-stage operational constraints. Furthermore, safety requirements can be formally incorporated through Constrained Markov Decision Processes (CMDPs), ensuring that long-term optimization does not violate critical operational or safety boundaries.

#### *Promising research direction for PHM: Agentic AI*

Looking ahead, an emerging trend is to integrate PHM with agentic AI systems, autonomous agents capable of continuous monitoring, reasoning, planning, and acting to manage asset health [346,347]. While traditional PHM models identify *whether a system is behaving anomalously* or *when the end of life will be reached*, Agentic AI addresses the Agency Gap by autonomously determining *how* to intervene. These agent-based frameworks leverage advances in RL and Large Language Models (LLMs) to achieve higher levels of autonomy and adaptivity in maintenance decision-making. In contrast to pre-defined or rule-based maintenance policies, agentic AI systems can learn optimal strategies online, reason over complex scenarios with domain knowledge, plan high-level maintenance tasks, and execute interventions in real time. This paradigm shift opens additional opportunities for scalable and adaptive health management in fleet-wide settings [346,348].

One key enabler of agentic PHM is the integration of LLMs with RL-based agents for decision making in maintenance settings. Recent work has demonstrated that LLMs can translate natural-language task goals into dense reward signals for RL agents [349]. For example, Xie et al. [349] use an LLM to automatically generate dense reward functions from natural-language task descriptions, substantially reducing the need for manual reward engineering in robotic manipulation and locomotion benchmarks. Although these are not maintenance-specific, the same paradigm could, in principle, be used to derive machine-readable reward functions from human-specified maintenance objectives such as minimizing unplanned downtime, avoiding unsafe operating regimes, or satisfying spare-part budget constraints. In a complementary line of work, Liang et al. [350] introduced a "Code as Policies" paradigm where an LLM generates executable code that composes perception outputs and control primitives into multi-step manipulation behaviors. This line of work illustrates how LLMs can map high-level instructions and domain abstractions into structured, executable decision logic. In an agentic PHM setting, analogous mechanisms could be used to generate or adapt multi-step maintenance workflows (e.g., inspection-diagnosis-repair sequences), which are then validated or refined by an RL or optimization layer rather than being hand-crafted by engineers. Early PHM-specific prototypes already explore parts of this space. Lukens et al. [351] present LLM-based PHM copilots that sit in the loop of existing maintenance recommendation workflows, where LLM agents process alerts and historical data to propose candidate maintenance actions while keeping human subject-matter experts in control. Kirubanandan [352] further combines LLM-generated diagnostic action sequences with an external causal inference module, constructing causal knowledge graphs from historical traces and using counterfactual reasoning to rank alternative intervention sequences. At the industrial scale, Getz et al. [353] demonstrate how LLMs are used to augment, summarize, and structure large corpora of maintenance work orders in an aerospace manufacturing plant, accelerating the extraction of recurring failure patterns and candidate interventions and paving the way for LLM-powered maintenance co-pilots. Together, these works suggest that combining the semantic understanding of LLMs with learning-based control or decision modules could support PHM agents that are better aligned with human maintenance objectives and capable of producing recommendations in a form that is directly interpretable by engineers.

Another promising direction of agentic AI in PHM is the increasing adoption of multi-agent reinforcement learning (MARL) for distributed asset management and fleet-level control. Real-world maintenance often involves coordinating decisions across multiple systems (e.g. a fleet of aircraft engines or wind turbines). MARL techniques allow a team of

agents to learn cooperative policies, achieving global optimization while respecting local conditions. For instance, Xia et al. [354] developed a multi-agent deep RL framework to collaboratively optimize maintenance schedules for an entire aero-engine fleet, outperforming independent single-agent strategies in cost reduction and reliability. In parallel, Sebastián et al. [355] propose a physics-informed MARL framework for distributed multi-robot coordination problems, such as formation control and cooperative navigation, that imposes a port-Hamiltonian structure on the agents' policies and leverages self-attention over the interaction graph to ensure scalable, energy-consistent control. From a broader critical-infrastructure perspective, Rokhforoz et al. [356] propose a semi-distributed MARL framework for optimal bidding of generation units in electricity markets, where each unit is modeled as an agent and a graph convolutional network encodes the power-grid topology into the agents' policies. By explicitly exploiting the network structure, their approach improves both profitability and generalization across different grid topologies compared to standard deep RL. Although the objective is economic rather than maintenance-focused, the combination of MARL with graph-based representations is highly relevant for PHM scenarios in networked infrastructures or IoT systems, such as power grids, pipeline networks, or interconnected production lines, where local decisions must account for system-wide couplings [6]. Beyond low-level control and coordination, agentic AI is also emerging at the orchestration layer. In smart manufacturing, Farahani et al. [348] propose a hybrid agentic AI and multi-agent framework for a prescriptive maintenance use case, where an LLM-based planner agent orchestrates workflows across perception, preprocessing, analytics, and optimization layers, coordinating specialized software agents in industrial settings. These developments suggest fleet-level PHM architectures where RL/MARL agents handle local asset- or subsystem-level decisions, orchestrated by higher-level LLM planners that reason over fleet-wide constraints, documentation, and human feedback.

While HRL appears earlier as a general learning strategy for improving exploration and sample efficiency, within agentic PHM it plays a distinct role as the hierarchical planning and decision backbone. Many PHM decision problems are inherently long-horizon and procedural: inspections, shutdowns, and component replacements must follow domain-specific workflows, safety constraints, and interdependent steps. Flat Markov decision process formulations often struggle to capture this structure, leading to myopic or hard-to-interpret policies. To address this, hierarchical frameworks decompose complex missions into sub-tasks and sequential plans. HRL provides the hierarchical decision backbone for agentic systems by enabling the decomposition of long-horizon tasks into simpler subproblems and the joint learning of high-level choices (e.g., which component or subsystem to inspect or repair next) and low-level skill policies [345], making it well-suited in principle for multi-stage maintenance scenarios where diagnostic, inspection, and repair actions must be coordinated over extended time horizons.

Finally, an essential attribute of agentic AI in PHM is continuous learning, the ability for maintenance agents to self-improve over time. Industrial environments are non-stationary: equipment degrades, operating regimes shift, and new failure modes appear. Approaches like meta-reinforcement learning enable agents to rapidly adapt to new task conditions using minimal additional experience [357]. In PHM settings, Cheng et al. [358] demonstrate that meta-RL can transfer maintenance strategies across different units in a fleet. Combined with digital-twin-based online learning pipelines, these adaptive agents can refine their predictive and prescriptive policies continuously as new sensor evidence accumulates.

In summary, the convergence of RL, LLMs, hierarchical planning, and multi-agent coordination is laying the groundwork for truly agentic PHM systems. Emerging works demonstrate the potential: LLM-integrated maintenance co-pilots [351,353], cooperative MARL frameworks [354,359], and adaptive meta-learned agents [358]. While many challenges remain, such as safety assurance, trust, and real-time compute constraints, the trajectory points toward PHM systems capa-

ble of autonomously managing asset health, coordinating maintenance actions, and learning continuously from experience.

## 5. Scaling beyond single systems

Physics-informed machine learning improves over purely data-driven methods by increasing model robustness and encouraging physically-consistent behavior. However, PHM methods still face deployment and scalability challenges, particularly in fleet-level scenarios.

When a model trained on data from a single unit is applied to other units within the same fleet, performance often degrades significantly [3]. Indeed, operational conditions can vary across assets, with many scenarios not captured in the original training data [360,361]. For example, in fault diagnosis, a well-trained deep model may underperform on a new machine simply due to unseen environmental or operational conditions. Similarly, if a machine's conditions evolve over time, the model may misinterpret these changes as faults if they deviate too much from the training distribution.

These challenges arise from the observational bias inherent in real-world data: models are rarely trained on datasets that capture the full range of operating conditions across a fleet, even though all units may share the same underlying physics [362]. This issue is further exacerbated by the limited availability of labeled data, as industrial datasets are often sparsely annotated due to the high cost of labeling.

While domain adaptation techniques aim to address distributional gaps between source and target domains, they typically assume access to sufficient unlabeled target data. However, in practice, PHM applications often involve small, imbalanced, and task-specific datasets with limited or no labels—conditions under which conventional learning methods perform poorly [363].

To overcome these challenges, fast adaptation and domain generalization have emerged as promising strategies for scaling PHM beyond single-system training. These techniques aim to enable quick model specialization or improve robustness across diverse conditions:

- **Fast Adaptation Methods (Section 5.1):** These techniques enable models to quickly specialize to new assets using only a small number of labeled examples, either by learning optimal starting parameters that can be fine-tuned with minimal data, or by leveraging in-context learning approaches that adapt without parameter updates.
- **Domain generalization methods (Section 5.2):** These approaches aim to learn invariant representations that remain stable across diverse operating conditions.

### 5.1. Fast adaptation methods

#### Introduction to fast adaptation methods

Physics-Informed Machine Learning (PIML) has advanced PHM by embedding physical principles into machine learning models, leading to better generalization and data efficiency, especially in safety-critical and data-sparse domains. However, real-world PHM applications often encounter rapidly changing conditions, new system configurations, and unexpected operational scenarios—that require not just robust, physics-consistent models but also the ability to adapt quickly.

Recent research has sought to address fast adaptation by developing adaptation and transfer learning mechanisms directly within PIML frameworks, enabling models to update or recalibrate in response to new data or shifts in system dynamics. For example, a recently proposed deep transfer operator learning framework enables fast and efficient learning of heterogeneous tasks despite substantial differences between source and target domains [371]. At the same time, fast adaptation methods from meta-learning and in-context learning provide complementary capabilities for purely data-driven PHM models, allowing them to learn from limited examples, exploit contextual information, and generalize across tasks and domains with minimal manual intervention. This is particularly valuable in fleet-scale PHM, where operational condi-

tions vary widely and where obtaining sufficient labeled data to retrain models across diverse operating scenarios is increasingly impractical at scale.

Fast adaptation methods build on knowledge acquired during an initial training phase across multiple related tasks or systems. This knowledge is embedded into the model's parameters or conditioning mechanisms and can be efficiently reused when adapting to new, unseen systems or operating conditions. The unifying objective of such approaches is to leverage prior experience to enable rapid task transfer under scarce supervision and shifting operating conditions, despite differences in formal assumptions across the different approaches. This problem is often framed within the paradigm of few-shot learning, where a model must adapt to a new task after seeing only a few labeled examples [372]. In the standard N-way K-shot setting, the model is required to classify inputs into N categories after observing only K labeled samples per class, with K typically ranging from one to ten [373]. Few-shot learning methods can be broadly grouped into two categories: those that explicitly update model parameters for each new task, and those that avoid weight updates by relying instead on prompting or conditioning mechanisms for fast adaptation. An important special case is zero-shot learning, corresponding to the  $K = 0$  limit of few-shot learning. Zero-shot learning is particularly attractive for PHM applications, as it allows deployment on unseen assets, operating conditions or fault types without the need for costly data collection or retraining. Instead, it exploits prior knowledge and structural similarities between tasks to generalize across domains [374–376].

A key distinction among fast adaptation methods lies in whether they require parameter updates during adaptation or operate in a parameter-free manner. Metric-based few-shot methods, such as Siamese networks [377], Prototypical Networks [378], and Relation Networks [379], fall into the latter category. These approaches learn an embedding space where samples from similar classes cluster closely together, enabling classification by comparing distances between embeddings rather than retraining the model.

Meta-learning, often described as "learning to learn," takes a different approach by training across a distribution of related tasks to discover a parameter initialization that is broadly useful. This initialization allows for rapid adaptation to new tasks or systems with only a few labeled examples [373,380,381]. Among the many meta-learning frameworks, Model-Agnostic Meta-Learning (MAML) [382] has gained particular traction in PHM research because of its architecture independence. MAML directly optimizes model parameters to find a general initialization, defined as a point in parameter space from which gradient-based updates quickly converge to task-specific optima, as illustrated in Fig. 10. The algorithm uses a two-level optimization process: an inner loop that adapts the model to individual tasks using a few examples, and an outer loop that updates the initial parameters based on how well the model performs across many different tasks. In PHM, this approach enables the model to learn a good starting point that generalizes across the fleet, so that for example, when deployed on a new asset, the model can quickly adapt using only limited data from that specific unit [383,384].

The category of parameter-free, prompt-conditioned fast adaptation methods keeps the model weights frozen and performs adaptation entirely at inference time [385]. Unlike few-shot learning, which typically fine-tunes model parameters using a small set of labeled examples, in-context learning (ICL) typically leverages large pre-trained foundation models by concatenating the support examples with the query directly in the input context. First identified in large language models [386], this "few-shot prompting" capability has since been demonstrated across various types of foundation models, including those designed for graph-structured data and multimodal scenarios [387,388]. In-context learning has proven effective for both classification and regression tasks using real and synthetic pre-training data, relying on tabular foundation models [367,389], and has been evaluated extensively in the PHM context [369]. Using a time generative adversarial network for feature extraction, tabular foundation models have been applied to slewing

**Table 11**  
Checklist of fast adaptation methods.

Applications in PHM	Type of tasks	Fault detection; Fault diagnosis; RUL prediction.
	Addressed challenges	Fast adaptation to a novel dataset or task.
Requirements	Prior knowledge	Consistent observational biases across small datasets.
	Data	Dataset with few samples (5-100); Multiple small data sets; Unlabeled data
	Type of bias	Observational bias.
	Assumptions	Observational biases can be learned by a single model, regardless of the task. This means a training task can be formulated at the meta-level, or a suitable pre-training task can be defined.
Summary	Advantages	Fast adaptation to novel tasks; Data efficiency; Robustness to distribution shifts; Integration of various data sources.
	Disadvantages	Computational inefficiency; Difficult task definition; Slow Adaptation in complex tasks; Limitations on task diversity
Resources	Repositories	<b>MAML</b> [364] <sup>a</sup> : original implementation in TensorFlow; <b>pytorch-maml</b> [364] <sup>b</sup> : PyTorch implementation; <b>Torchmeta</b> [365] <sup>c</sup> : a collection of extensions and data-loaders for few-shot learning & meta-learning; <b>learn2learn</b> [366] <sup>d</sup> : a software library for meta-learning research ensuring fast prototyping and reproducibility; <b>higher</b> [367] <sup>e</sup> : a library for higher-order optimization; <b>TabPFN</b> [368] <sup>f</sup> : a tabular foundation model for predictions on small data; <b>TabPHM</b> [369]: an evaluation framework for tabular foundation models tailored to PHM applications; <b>AnomalyGPT</b> [370] <sup>g</sup> : A large vision-language model for industrial anomaly detection.
	Datasets	<i>Fast Adaptation Methods</i> are not tied to specific datasets.

<sup>a</sup> <https://github.com/cbfinn/maml>

<sup>b</sup> <https://github.com/tristandeleu/pytorch-maml>

<sup>c</sup> <https://github.com/tristandeleu/pytorch-meta>

<sup>d</sup> <https://learn2learn.net/>

<sup>e</sup> <https://github.com/facebookresearch/higher>

<sup>f</sup> <https://github.com/PriorLabs/TabPFN>

<sup>g</sup> <https://anomalygpt.github.io/>

bearing fault diagnosis in combination with other deep learning models [390].

#### Application of fast adaptation methods in PHM

Fast adaptation is a central objective in PHM and is often studied through the lens of few-shot learning or related meta-learning paradigms. In PHM, **few-shot learning** has been explored extensively to tackle label-scarce classification problems, particularly in fault diagnosis [381]. Early applications of few-shot learning in PHM primarily built on established methods, combining existing metric-based few-shot learning methods, such as Siamese and prototypical networks, with neural architectures that were already proven effective in the field, particularly convolutional neural networks. For example, Yang et al. [391] demonstrate how a few-shot learning method can be applied in bearing fault diagnosis, employing Siamese convolutional neural networks, while Zhao et al. [392] propose a convolutional prototype network for chemical process fault diagnosis. More recent work in metric-based few-shot learning has begun to integrate transformer-inspired architectures, especially attention mechanisms. For instance, Hu et al. [393] propose a relational convolutional block attention network for bearing fault diagnosis under few-shot conditions, specifically targeting the imbalance between abundant artificial damage data and limited real-world damage data. Building on this direction, Wang et al. [394] proposed an enhanced transformer model that leverages the same attention mechanism but is trained with a novel asymmetric loss function to improve robustness against label noise in few-shot bearing fault diagnosis under variable speed conditions. Complementing these architecture-focused advances, recent research has also revisited the design of distance measures for metric few-shot learning. Costa et al. [395], for instance, employ a pre-trained variational autoencoder and introduced a novel Kolmogorov-complexity-based distance in the latent space, enabling system health estimation with minimal labeled data based on most similar k-neighbors. In parallel, other works have shifted focus from distance metrics to the learning paradigm itself, investigating how to decouple feature extrac-

tion from episodic training and integrate meta-learning with multi task learning. Wang et al. [396] address this by first training a feature extractor through supervised classification on source-domain labels, then freezing it and performing episodic metric learning in the resulting feature space. This two-stage process improves knowledge transfer and yields higher accuracy for bearing and gearbox fault diagnosis under limited data conditions. Building on this idea, Lei et al. [397] demonstrate that feature extraction can also be enhanced through semantic self-supervised pretraining. They propose a prior knowledge-embedded meta-transfer learning framework that jointly optimizes multi-task metric meta-learning and adaptive information fusion, achieving robust performance in bearing fault diagnosis under scarce-label settings.

While metric-based few-shot learning primarily targets classification tasks, **meta-learning** offers a more versatile framework for fast adaptation in PHM [398]. Due to its flexibility and model-agnostic nature, MAML has become one of the most widely adopted meta-learning approaches in the PHM literature. Applications span from the health monitoring of rotary machine components, such as bearings, to machining and manufacturing processes, as well as turbofan engine prognostics. For example, Ding et al. [399] integrate MAML into a meta learning framework for bearing prognostics and fault diagnostics, while Yang et al. [400] employ MAML with adaptive threshold networks to improve cross-domain fault diagnosis. Building on these classification-oriented studies, large body of work extends MAML to regression tasks, most notably for remaining useful life estimation in turbofan engines [401–403], as well as for tool wear prediction in machining processes [404].

Recently, several extensions of MAML have been proposed to better address the specific challenges of PHM tasks. For example, Wang et al. [403] extend MAML to a Bayesian framework, enabling probabilistic modeling of multiple failure modes and more reliable estimation of remaining useful life [403]. Building on this line of work, Yang et al. [405] integrate MAML with conditional normalized flow to infer richer posterior distributions in the kernel features, thereby extracting more discriminative information from engine and bearing degradation data. In a related direction, Ding et al. [406] combine MAML with

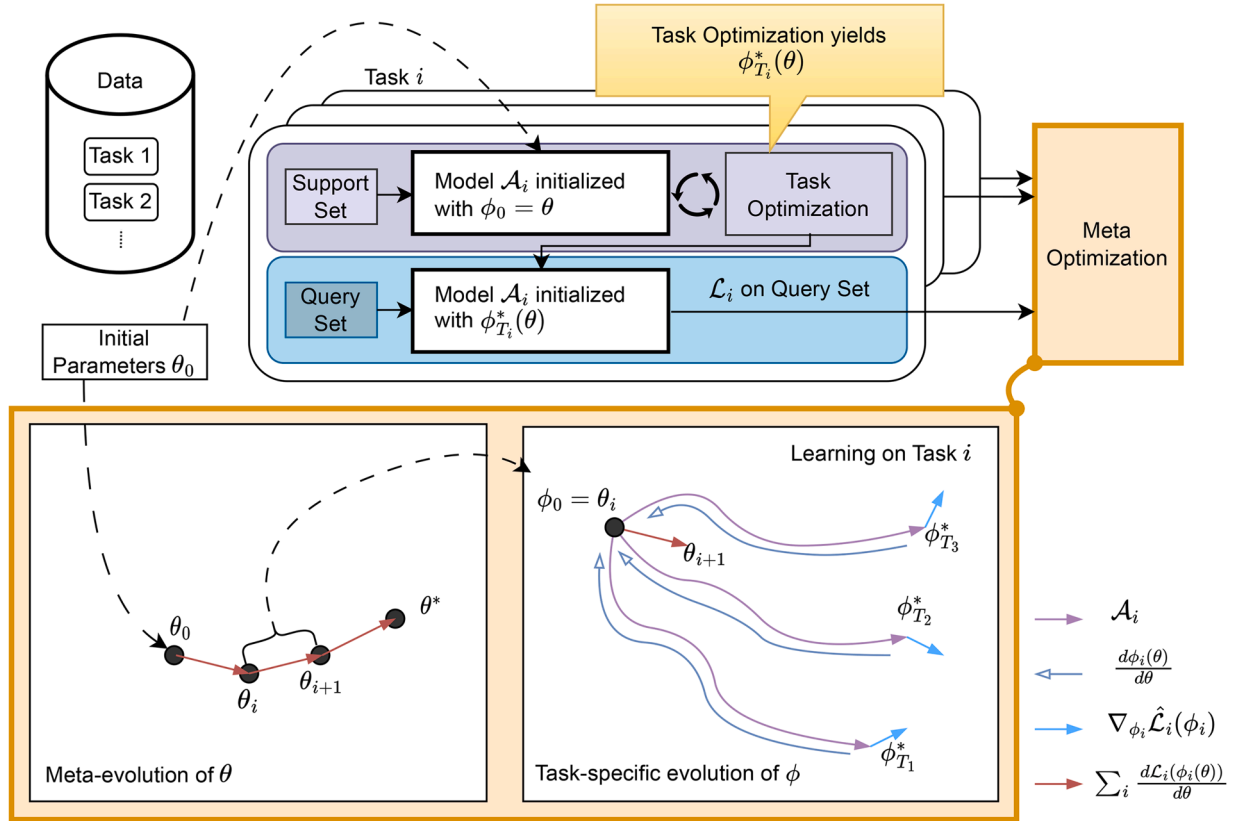


Fig. 10. Model-Agnostic Meta-Learning (MAML). For each task, the model starts from shared meta-parameters  $\theta_0$ , adapts on a support set to obtain task-specific parameters, and then uses the query set to update  $\theta$  for better generalization. The bottom-left inset shows meta-parameters converging toward the optimal initialization  $\theta^*$ , while the bottom-right inset illustrates task-specific adaptation trajectories.

pseudo labeling to enable few-shot prognostics when large amounts of unlabeled historical data are available.

Beyond these extensions of MAML, meta-learning has also been integrated with GNNs for remaining useful life prediction under variable operating conditions by constructing a spatio-temporal graph that integrates structural data with multi-source, multi-channel signal, demonstrating the applicability of meta-learning to non-euclidean domains [407]. In addition to graph-based approaches, physics-informed meta learning has also been explored. Li et al. [408] propose a framework that integrates empirical tool wear models into data-driven modeling, allowing the meta-learning process to adapt quickly across wear stages while remaining consistent with physical laws. Similarly, Majumdar et al. [409] propose a hypernetwork-based physics-informed neural network model for real-time monitoring of an industrial heat exchanger. While a large body of literature has focused on MAML-based strategies, recent results demonstrate that metric-based few-shot approaches remain competitive, and in fact outperform MAML on two bearing datasets using recent transformer architecture [394].

Zero-shot methodology has recently gained traction in PHM as a way to detect previously unseen fault types without requiring labeled training data. Approaches in this direction span several complementary strategies. Some integrate physical knowledge with generative models, embedding physics-based priors for zero-sample diagnosis [410]. Others leverage semantic descriptions of faults, where bi-linear compatibility functions map unseen to seen classes [374].

Building on semantic spaces, Ma et al. [411] diagnose compound faults from single-fault training data by projecting vibration features into a descriptive semantic space. Chen et al. [412] exploit information granularity, predicting pyramid-structured attributes and matching them to expert-defined descriptions. Xu et al. [375] cluster representations in semantic space to infer category centroids for unseen fault

detection, while Cai et al. [413] combine out-of-distribution detection with clustering and semantic attribute projections to refine classification of both seen and unseen faults.

Although no dominant paradigm has yet emerged, this growing body of work demonstrates the promise of zero-shot learning for generalizing fault diagnosis to entirely new failure modes.

In parallel to few-shot and meta-learning, **in-context learning** has so far seen limited adaptation in PHM, though it shows promise in related anomaly detection tasks. For instance, Zhu and Pang [414] proposed an in-context residual learning framework that detects anomalies across diverse datasets, spanning industrial, medical, and semantic datasets, without requiring additional training, demonstrating the adaptability of ICL to heterogeneous settings. Building on this idea, Gu et al. [370] employed a large vision-language model to generate textual anomaly scores in few-shot in-context scenarios, aiming to enhance transparency and interpretability in industrial anomaly detection. These early efforts suggest that ICL could offer a new paradigm for PHM by enabling rapid, training-free adaptation with improved interpretability. A summary of the application contexts for fast adaptation methods is provided in Table 11.

#### Advantages and limitations of fast adaptation methods

Parameter-updating methods such as MAML learn parameter initializations that capture shared patterns across operating conditions or fleet units, enabling rapid adaptation to new configurations or assets with minimal labeled data and few gradient steps [415]. This significantly reduces both calibration time and the reliance on extensive fault libraries for each unit. However, the bi-level optimization introduces substantial computational overhead from backpropagating through inner and outer loops, which can be prohibitive in industrial applica-

tions [382,416]. In addition, MAML performance remains dependent on the underlying model architecture. These methods are also sensitive to domain shift, when deployment conditions differ from those seen during meta-training, the learned initialization may adapt poorly or produce miscalibrated predictions [417]. Additionally, changes in sensor modalities or data structures typically require retraining, since the meta-learned parameters are architecture-dependent [418]. Compared to MAML, metric-based few-shot methods have the advantage of avoiding inner-loop gradient updates at test time, which makes them faster, more stable, and less sensitive to optimization hyperparameters, as inference typically reduces to embedding and nearest-neighbor comparison rather than meta-optimization. Nevertheless, these methods can struggle to generalize when new classes or tasks significantly differ in distribution from the training data, since the learned metric or embedding space may fail to capture meaningful similarities and the resulting metric can become unreliable in unseen domains.

Parameter-free approaches like in-context learning enable adaptation by providing a few labeled examples directly within the input, allowing the model to infer task patterns without any parameter updates. This offers significant flexibility since the same model can handle diverse tasks and data structures simply by changing the provided examples. However, these methods rely on attention mechanisms that scale quadratically with sequence length, creating substantial computational overhead that limits both the length of time series that can be processed [419] and the number of in-context examples that can be provided. Performance depends heavily on carefully selecting, ordering, and curating the in-context examples [172,420].

#### *Unexplored and emerging applications of fast adaptation methods in PHM*

Building on recent advances in meta-learning advances opens up several promising research directions for PHM. Meta-reinforcement learning, for instance, shows how agents can acquire learning priors that enable rapid adaptation to new tasks from few interactions [421]. Applied to PHM, this could yield control policies that adjust on the fly to new machines, fault modes, or operating conditions, alleviating the need to retrain from scratch for each scenario and reducing the computational burden of RL applications (Section 4).

Complementing this, physics-informed meta-learning shows strong potential for generating surrogate models that generalize across PDE-driven tasks and adapt to new system parameters with minimal data [422,423]. For PHM, such approaches could enable fast-calibrating digital twins that adapt to new assets with limited sensor data while respecting known degradation physics.

A third direction is unsupervised meta-learning, which generates diverse training tasks from unlabeled data through augmentations and mixing strategies to train models for transferable adaptation patterns [424]. Unlike foundation models where in-context learning emerges incidentally from general pretraining, this approach explicitly optimizes for task adaptation. In PHM, this could allow models to learn adaptation strategies directly from abundant unlabeled operational data and later specialize to new fault modes or asset configurations without extensive labeled fault libraries.

Finally, foundation models present a complementary approach. General-purpose language models have been adapted for equipment monitoring in industrial cyber-physical systems [425] and wind turbines [426], with specialized variants developed for mechanical fault diagnosis of bearings [427]. Recent advances beyond language models show that time-series foundation models can be reprogrammed by keeping language-model backbones frozen while adding lightweight adapters for cross-domain transfer [428], or by treating series as text tokens to enable zero-shot forecasting [429]; unified architectures further handle forecasting, classification, and anomaly detection within a single backbone [430], with multimodal grounding and federated variants improving instruction-following over visualized traces and privacy-preserving cross-domain transfer [431,432]. These capabilities have been success-

fully transferred to fault detection applications [433,434]. However, currently available publicly released foundation models still lack domain specificity. Several works therefore investigate the development of customized foundation models that integrate domain-specific priors of large-scale problem areas such as the monitoring of wind turbines [426] or electric power grids [435].

#### *5.2. Domain generalization: Learning to generalize to unseen domains*

##### *Introduction to domain generalization*

A fundamental challenge in PHM is domain shift, which occurs when the data distribution in the operational environment differs from the training dataset. This distributional discrepancy arises from several factors, including variations in operating and environmental parameters [445] such as load, speed, and temperature, and humidity, as well as differences in system configurations such as sensor types and placements, calibration settings, or overall system design. Furthermore, temporal factors like system degradation and seasonal changes can introduce gradual data shifts [446]. Consequently, this mismatch between source and target domains often leads to a significant deterioration in model performance, limiting the generalizability of machine learning solutions to new conditions [447,448].

To address the challenge of domain shifts, several learning paradigms have been explored. While fast adaptation methods (Section 5.1) such as few-shot learning [449], meta-learning [450], in-context learning [170], and domain adaptation (DA) [448] are effective, they typically rely on the availability of samples from the target domain. However, this assumption may not hold in many real-world scenarios, such as fault diagnosis and remaining useful life prediction. Domain generalization (DG) [437,451], in contrast, offers a more promising approach. By leveraging knowledge from several source domains, DG techniques aim to develop models that exhibit improved generalization capabilities on target domains not encountered during training. An illustration of the differences between domain adaptation and generalization is shown in Fig. 11.

DG can be generally categorized into three main approaches [437]: data manipulation, representation learning, and learning strategies. Data manipulation techniques enhance generalization by increasing the diversity of training data through augmentation [452,453]. Representation learning methods seek to acquire domain-invariant representations using domain-adversarial neural networks [454,455], explicit feature distribution alignment [456], and instance normalization [457]. Additionally, learning strategies such as meta-learning [458], gradient operation [459], and self-supervised learning [460] have been explored to improve generalization performance. For example, [459] forces the network to activate features that correlate with labels, and [460] designs multimodal self-supervised pretext tasks such as masked cross-modal translation and multimodal Jigsaw puzzles to improve generalization. **Prior Knowledge-based Domain Generalization** Beyond purely data-driven strategies, domain generalization can also benefit from incorporating prior knowledge that captures invariances shared across domains. Such priors constrain the learning process toward features that are less sensitive to distributional shifts, thereby improving robustness and transferability.

Variations in operating conditions, environmental factors, and system configurations described above can alter the statistical properties of the measured signals, often affecting their frequency content. Since certain frequency components remain relatively stable across domains, they can serve as robust priors for domain generalization. Leveraging such frequency-domain priors allows models to focus on features that are invariant to domain shifts.

For example, the Fourier phase of an image, which encodes high-level semantic structures, is robust to common domain shifts [461]. Similarly, deep neural networks exhibit an inherent bias towards learning certain frequency components during training [462]. Incorporating these frequency-based priors into the learning objective can

**Table 12**  
Checklist of prior knowledge-based domain generalization.

Applications in PHM	Type of tasks	Fault detection; RUL prediction.
	Addressed challenges	Poor generalization to different machines or conditions.
Requirements	Prior knowledge	Frequency components in neural networks and Fourier-based analysis.
	Data	Labeled source domain data.
	Type of bias	Frequency domain prior of data and neural networks.
	Assumptions	The labeled source domain data is available while the target domain data is unavailable.
Summary	Advantages	Improved generalizability with incorporation of Frequency domain information.
	Disadvantages	Difficulty of capturing all possible domain variations and the risk of overfitting to source domains.
Resources	Repositories	DomainBed [436], Transfer Learning Library [437], DG for Fault Diagnosis [438]
	Datasets	XJTU Dataset: acceleration signals from bearing [49], PU Dataset: artificially damaged and life-accelerated bearings [439], JNU Dataset: bearings in a centrifugal fan system [440], SCP Dataset: self-priming centrifugal pump [441], MIMI: sound of industrial machines under normal and anomalous operating conditions [442], PACS: object recognition benchmark including four domains (photo, art-painting, cartoon, sketch) [443], OfficeHome: object recognition benchmark from four domains (Art, Clipart, Product, Real-World) [444].

guide models toward more transferable features, thereby improving robustness and generalization.

A common technique involves manipulating the Fourier components directly, where phase information (structural content) is preserved while the amplitude information (domain-specific characteristics) is swapped with that of other domains [461]. This strategy retains core structural information while mitigating domain-specific variations. To further refine this, some methods introduce a soft-thresholding function to remove small, potentially noisy frequency components that might be domain-specific [463].

Beyond static Fourier manipulation, recent methods combine frequency priors with physically grounded augmentation strategies. For instance [464] perturbs the Fourier amplitude in proportion to the original energy distribution, ensuring that the perturbations remain physically plausible while introducing controlled variability. This targeted spectral modification exposes the model to a broader range of frequency configurations, encouraging the extraction of structural cues that remain stable across domains. Similarly, PhysAug [465] integrates physical guidance with frequency-domain augmentation for single-domain generalized object detection. By combining physically motivated perturbations with spectral manipulations, PhysAug improves model robustness without introducing unrealistic artifacts.

In addition to spectral invariance, prior knowledge can also take the form of structural and causal relationships derived from the underlying physics of the system. For instance [466] encodes causal dependencies between physical variables and measurement signals, integrating them into the learning process via physics-based constraints. By constraining the learned features to adhere to known physical mechanisms, the model is encouraged to focus on domain-invariant relationships, improving generalization to operating conditions never encountered during training.

Finally, the inherent frequency bias of deep neural networks can itself be leveraged or corrected. Such models tend to learn lower-frequency components first, which can create a dependency on "frequency shortcuts" patterns that are highly correlated with the source domain but do not generalize well [467]. Recognizing this, Lin *et al.* [462] introduced deep frequency filtering, which modulates frequencies in the latent space. This method uses an attention mechanism in the frequency domain to amplify transferable features and suppress non-transferable ones.

#### Application of domain generalization in PHM

In many practical PHM applications, obtaining representative labeled data from all possible future operating conditions or equipment

configurations is either prohibitively expensive or simply not feasible. In such cases, DG becomes a more critical and relevant task than DA, as it does not rely on access to target domain data during the training phase. In PHM, the successful application of DG enhances the reliability of key predictive tasks, such as RUL estimation and fault diagnosis, across diverse and previously unseen equipment. This improved predictive accuracy directly enables more efficient resource allocation. For example, by facilitating just-in-time maintenance scheduling based on reliable RUL predictions, organizations can avoid both premature component replacement and costly unexpected failures. This targeted approach optimizes the deployment of maintenance personnel and spare parts, ultimately maximizing equipment availability and reducing operational costs. The following sections review how DG approaches have been applied to the two labeled PHM tasks: fault diagnosis and remaining useful life prediction.

Domain generalization plays a crucial role in building fault diagnostics and prognostic models that are robust to unseen operating conditions. Key strategies include data manipulation, representation learning, and advanced learning strategies [468]. Data manipulation techniques enhance training data diversity [469]; for instance, the domain augmentation generalization network employs multi-source augmentation and adversarial training to generate varied data, boosting the model's robustness to out-of-distribution samples [470]. Representation learning aims to extract domain-invariant features [471,472]. A common strategy is to combine Domain Adversarial Neural Network for implicit feature alignment with explicit Maximum Mean Discrepancy minimization to reduce distribution gaps, an approach that has been successfully applied to diagnose rock drill faults under variable conditions [473]. Finally, learning strategies like meta-learning directly optimize for generalization. The Meta-GENE framework [474], for example, uses a model-agnostic meta-learning procedure with gradient alignment to learn a domain-invariant optimization policy, achieving strong performance on gearbox fault diagnosis tasks.

Applying domain generalization to RUL prediction is essential for creating prognostic models that can adapt to new assets or changing operational profiles, a task where target run-to-failure data is often not available [475]. RUL prediction is a regression task and is more difficult than classification because it requires precise, continuous predictions that are highly sensitive to distribution shifts. Additionally, small changes in input can lead to large output errors, making robust generalization across domains more challenging than in classification. Data manipulation strategies generate synthetic degradation trajectories to train more generalizable regressors [476]. The Adversarial Out-Domain Augmentation [477] framework exemplifies this by training a generator to produce diverse pseudo-domains of degradation signals, which are then

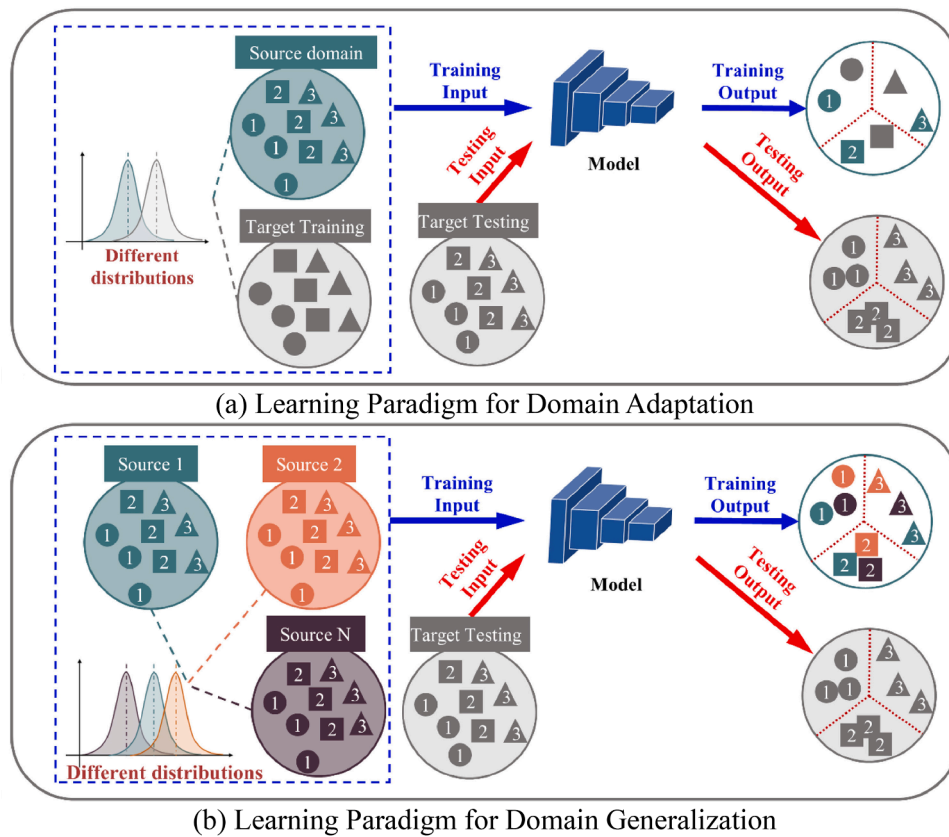


Fig. 11. Differences between learning paradigms for domain adaptation and domain generalization [438].

used to train a robust RUL predictor. Representation learning methods focus on creating a domain-agnostic health index [478]. This is often achieved by integrating recurrent architectures like Bidirectional Gated Recurrent Units (BGRU) [479] with a DANN, where the BGRU captures temporal degradation patterns and the DANN aligns feature distributions across different working conditions. Advanced learning strategies are also employed, particularly for small-data scenarios. For instance, a meta-domain generalization method for tool wear prediction [480] uses a phased approach of data splitting, meta-pre-training, and fine-tuning to effectively generalize from limited source data. A summary of the application contexts for domain generalization is provided in Table 12.

#### Advantages and limitations of domain generalization in PHM

Domain generalization offers significant advantages for the practical application of PHM in industrial settings. A primary benefit is the reduced dependency on labeled data from every potential target domain. In many PHM applications, acquiring comprehensive labeled datasets that encompass all operating conditions, equipment variations, and environmental factors is often prohibitively expensive, time-consuming, or impractical [438]. DG mitigates this challenge by enabling models trained on a limited set of source domains to generalize effectively to new, unseen target domains, thereby reducing the need for target-specific labels. Additionally, DG does not rely on access to target domain samples, thereby addressing a key limitation of DA methods.

Moreover, models developed with DG techniques demonstrate enhanced robustness and reliability when deployed in real-world environments characterized by distribution shifts. By design, these models are more adept at managing the inherent variability of operational data, resulting in more consistent and accurate predictions of faults and RUL across diverse conditions [481]. This enhanced reliability is essential for building trust in PHM systems and supporting their adoption in safety-critical industrial and infrastructure applications.

Despite these benefits, DG faces notable limitations. It fundamentally relies on the assumption that the source domains capture enough variability to represent potential target domains. When novel or unforeseen operational shifts occur in the target environment, they may lie outside the scope of the source data, leading to degraded performance. In addition, strong generalization typically requires a diverse and carefully curated training dataset, which can be difficult to obtain in many industrial PHM settings.

#### Unexplored or emerging applications of domain generalization in PHM

In PHM, domain generalization is an emerging approach to enhance fault diagnostics and prognostics across diverse operational conditions. However, most existing DG research in PHM primarily focuses on unimodal data, while the real world is multimodal. Effectively integrating information from multiple modalities is crucial for improving diagnostics and prognostics performance [482,483]. One major barrier to multimodal DG in PHM is the lack of public multimodal datasets and benchmarks. For example, the accelerometer data is the predominant modality in fault diagnosis, while other modalities may also help, such as images and audio data, as motivated in Section 4. To advance research in this area, it is essential to identify use cases where multiple modalities are critical and develop standardized benchmarks tailored to those applications. For example, in the railway monitoring system, vision modality is a good complement to time-series modality for arcing detection. An additional limitation is that most existing DG approaches for PHM operate under the closed-set assumption, which presumes identical label spaces across domains. However, in real-world applications, target domains often contain unknown faults, making it essential to detect and handle them effectively. Techniques developed from the out-of-distribution detection [484,485] and anomaly detection [486] domains could be leveraged and adapted to facilitate unknown class detection. Finally, exploring DG techniques for fault detection is also

promising, as it addresses a fully unsupervised and considerably more challenging setting.

## 6. Conclusion

Prognostics and Health Management has evolved into a critical discipline for ensuring the safety, efficiency, and resilience of complex engineered systems. Modern PHM extends well beyond traditional tasks such as fault detection, diagnostics and prognostics to encompass system health management, health-aware decision-making, real-time control, and fleet-level asset operation and maintenance optimization. Meeting these demands requires algorithms that are accurate and scalable, interpretable, robust to uncertainty, and consistent with physical system behavior.

While machine learning has enabled major advances in PHM by modeling complex dynamics and integrating multimodal data, purely data-driven models often struggle under distribution shifts, perform poorly when data is limited or noisy, and often lack physical consistency. These limitations are particularly critical in domains where degradation processes are gradual, only partially observable, and influenced by external factors not directly measured. To address these challenges, Physics-Informed Machine Learning has emerged as a promising paradigm that integrates domain knowledge and physical constraints into the learning process, improving model reliability and trustworthiness.

In this second part of our review on Physics-Informed Machine Learning for PHM, we have focused on the remaining two biases that complement the inductive biases reviewed in Part I: **learning biases** and **observational biases**. These two biases shape how machine learning models are trained and exposed to data in ways that embed physical knowledge, improve robustness, and enhance trustworthiness, especially in industrial settings where data is sparse, labels are limited, and degradation evolves under complex, partially observable conditions.

We examined how learning biases guide the optimization process by embedding physical principles into loss functions, training objectives, or regularization schemes. Physics-Informed Neural Networks and degradation-informed learning methods serve as prototypical examples, ensuring that predictions are consistent with known dynamics or general physical properties of degradation processes such as monotonicity and irreversibility. In parallel, observational biases enhance model robustness by curating, enriching, or reconstructing datasets that more faithfully represent system behavior. Techniques such as virtual sensing, overcoming the simulation to real domain gap of physics-based simulations, physics-informed data augmentation, and multimodal fusion help bridge gaps in observability, enabling models to generalize beyond narrowly defined training distributions.

We also highlighted how to transition from prediction to action in PHM by leveraging reinforcement learning and health-aware control, enabling intelligent decision-making grounded in physical understanding. In this context, RL closes the loop from physics to machine learning and back to the physical world, where model-driven decisions directly affect system behavior. This shift transforms PHM from a passive monitoring task into an interactive, decision-oriented discipline, where learned policies adaptively manage reliability, risk, and performance in dynamic environments.

Despite this progress, several open challenges remain. A major limitation, as already outlined in Part I, is the lack of systematic benchmarking and validation in realistic industrial contexts. Methods are typically evaluated on synthetic or clean and controlled laboratory datasets, making it difficult to assess their robustness, transferability.

Degradation-informed approaches provide valuable flexibility, but their simplifying assumptions remain limiting, particularly in disentangling degradation from operational dynamics. Observational bias techniques such as leveraging simulated data or deploying virtual sensors carry inherent limitations. Synthetic trajectories generated from physics-based simulations often fail to reflect the complexity and variability of real-world degradation, hence leading to a simulation-to-real

gap. Virtual sensing strategies that infer hidden system states from indirect sensor measurements introduce modeling and identifiability uncertainties, as sensor drift can lead to error accumulation that propagates misrepresentations of system health.

Looking ahead, several promising research directions emerge. One is the development of adaptive training strategies that adjust the influence of physical constraints based on model confidence or data coverage. Another is advancing physics-aware generative models that synthesize realistic and physically plausible degradation trajectories. In RL, there is increasing potential in embedding physical structure at different levels, including the reward function, policy architectures, world models, and value functions, enabling physics-informed and sample-efficient decision-making. Hybrid frameworks that combine partial physical knowledge, data-driven inference, and symbolic reasoning are particularly promising for scenarios where physics is implicit or only partially known. Finally, scaling to fleet-level deployment will require generalization across diverse systems and operating conditions, motivating new research in meta-learning, in-context adaptation, and domain generalization under physical constraints.

In summary, this second part of our review expands the PIML framework to include not only inductive biases but also learning and observational biases, thus completing the picture of how physical knowledge can be infused into data-driven PHM systems. By embedding physics throughout the modeling pipeline (from data to training, to decision-making) PIML offers a powerful paradigm for building scalable, reliable, and physically consistent PHM solutions that meet the demands of real-world deployment in complex, safety-critical systems.

## CRedit authorship contribution statement

**Olga Fink:** Writing – review & editing, Writing – original draft, Supervision, Resources, Project administration, Investigation, Funding acquisition, Conceptualization; **Ismail Nejjar:** Writing – review & editing, Writing – original draft, Visualization, Methodology, Investigation, Conceptualization; **Vinay Sharma:** Writing – review & editing, Writing – original draft, Methodology, Investigation, Conceptualization; **Keivan Faghhi Niresi:** Writing – review & editing, Writing – original draft, Methodology, Investigation; **Han Sun:** Writing – review & editing, Writing – original draft, Methodology, Investigation; **Hao Dong:** Writing – review & editing, Writing – original draft, Methodology, Investigation; **Chenghao Xu:** Writing – review & editing, Writing – original draft, Methodology, Investigation; **Amaury Wei:** Writing – review & editing, Writing – original draft, Methodology, Investigation; **Arthur Bizzi:** Writing – review & editing, Writing – original draft, Methodology, Investigation; **Raffael Theiler:** Writing – review & editing, Writing – original draft, Methodology, Investigation; **Yuan Tian:** Writing – review & editing, Writing – original draft, Methodology, Investigation; **Leandro Von Krannichfeldt:** Writing – review & editing, Writing – original draft, Methodology, Investigation; **Zhan Ma:** Writing – review & editing, Writing – original draft, Methodology, Investigation; **Sergei Garmayev:** Writing – review & editing, Writing – original draft, Methodology, Investigation; **Zepeng Zhang:** Writing – review & editing, Writing – original draft, Methodology, Investigation; **Mengjie Zhao:** Writing – review & editing, Writing – original draft, Methodology, Investigation; **Kevin Steiner:** Writing – review & editing, Writing – original draft, Investigation; **Yusuf Kesmen:** Writing – review & editing, Writing – original draft, Investigation.

## Acknowledgment of AI assistance in manuscript preparation

During the preparation of this work, the authors used ChatGPT to assist with refining and correcting the text. After using this tool, the authors carefully reviewed and edited the content as needed and take full responsibility for the content of this publication.

## Data availability

No data was used for the research described in the article.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgement

This work has been supported by the Swiss National Science Foundation (SNSF) Grant 200021\_200461.

## References

- [1] Fink O, Wang Q, Svensen M, Dersin P, Lee W-J, Ducoffe M. Potential, challenges and future directions for deep learning in prognostics and health management applications. *Eng Appl Artif Intell* 2020;92:103678.
- [2] Arias Chao M, Kulkarni C, Goebel K, Fink O. Hybrid deep fault detection and isolation: combining deep neural networks and system performance models. *Int J Prognosti Health Manage* 2019;10(4):033.
- [3] Wang Q, Michau G, Fink O. Domain adaptive transfer learning for fault diagnosis. In: 2019 Prognostics and system health management conference (PHM-Paris). IEEE; 2019, p. 279–85.
- [4] Chao MA, Kulkarni C, Goebel K, Fink O. Fusing physics-based and deep learning models for prognostics. *Reliab Eng Syst Saf* 2022;217:107961.
- [5] Karniadakis GE, Kevrekidis IG, Lu L, Perdikaris P, Wang S, Yang L. Physics-informed machine learning. *Nature Rev Phys* 2021;3(6):422–40. Publisher: Springer Nature. <https://doi.org/10.1038/s42254-021-00314-5>
- [6] Fink O, Sharma V, Nejjar I, Von Krannichfeldt L, Garmaev S, Zhang Z, et al. From physics to machine learning and back: Part I-learning with inductive biases in PHM. *Reliab Eng Syst Saf* 2026 ;271:112213. <https://doi.org/10.1016/j.res.2026.112213>
- [7] Vafa K, Chang PG, Rambachan A, Mullainathan S. What has a foundation model found? inductive bias reveals world models. In: Forty-second international conference on machine learning. 2025.
- [8] Lu L, Meng X, Mao Z, Karniadakis GE. DeepXDE: a deep learning library for solving differential equations. *SIAM Rev* 2021;63(1):208–28. <https://doi.org/10.1137/19M1274067>
- [9] Zubov K, McCarthy Z, Ma Y, Calisto F, Pagliarino V, Azeoglio S, et al. NeuralPDE: automating physics-informed neural networks (PINNs) with error approximations. 2021, arXiv preprint arXiv:210709443 .
- [10] Hao Z, Yao J, Su C, Su H, Wang Z, Lu F, et al. Pinnacle: a comprehensive benchmark of physics-informed neural networks for solving pdes. In: Advances in neural information processing systems. 2024, Datasets and Benchmarks Track.
- [11] Takamoto M, Praditia T, Leiteritz R, MacKinlay D, Alesiani F, Pflüger D, et al. Pdebench: an extensive benchmark for scientific machine learning. In: Advances in neural information processing systems. 2022, Datasets and Benchmarks Track.
- [12] Raissi M, Perdikaris P, Karniadakis GE. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J Comput Phys* 2019;378:686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>
- [13] Ameya DJ, Em Karniadakis G. Extended physics-informed neural networks (XPINNs): a generalized space-time domain decomposition based deep learning framework for nonlinear partial differential equations. *Commun Comput Phys* 2020;28(5):2002–41.
- [14] Yang L, Meng X, Karniadakis GE. B-PINNs: bayesian physics-informed neural networks for forward and inverse PDE problems with noisy data. *J Comput Phys* 2021;425:109913. <https://doi.org/10.1016/j.jcp.2020.109913>
- [15] Ramirez I, Pino J, Pardo D, Sanz M, Del Río L, Ortiz A, et al. Residual-based attention physics-informed neural networks for spatio-temporal ageing assessment of transformers operated in renewable power plants. *Eng Appl Artif Intell* 2025;139:109556.
- [16] Sholokhov A, Liu Y, Mansour H, Nabi S. Physics-informed neural ODE (PIN-ODE): embedding physics into models using collocation points. *Sci Rep* 2023;13(1):10166.
- [17] Mattheakis M, Sondak D, Dogra AS, Protopapas P. Hamiltonian neural networks for solving equations of motion. *Phys Rev E* 2022;105(6):065305.
- [18] Zhang Z-Y, Zhang H, Zhang L-S, Guo L-L. Enforcing continuous symmetries in physics-informed neural network for solving forward and inverse problems of partial differential equations. *J Comput Phys* 2023;492:112415.
- [19] de Beaulieu MH, Jha MS, Garnier H, Cerbah F. Remaining useful life prediction based on physics-informed data augmentation. *Reliab Eng Syst Saf* 2024;252:110451.
- [20] Sun B, Xue M, Su M. Damage detection of wind turbine blades via physics-informed neural networks and microphone array. *Energy* 2025;136859.
- [21] Shukla K, Di Leoni PC, Blackshire J, Sparkman D, Karniadakis GE. Physics-informed neural network for ultrasound nondestructive quantification of surface breaking cracks. *J Nondestr Eval* 2020;39(3):61.
- [22] Zheng B, Li T, Qi H, Gao L, Liu X, Yuan L. Physics-informed machine learning model for computational fracture of quasi-brittle materials without labelled data. *Int J Mech Sci* 2022;223:107282. <https://doi.org/10.1016/j.ijmecsci.2022.107282>
- [23] Wen P, Ye Z-S, Li Y, Chen S, Xie P, Zhao S. Physics-informed neural networks for prognostics and health management of lithium-ion batteries. *IEEE Trans Intell Veh* 2023;9:2276–89.
- [24] Wang F, Zhai Z, Zhao Z, Di Y, Chen X. Physics-informed neural network for lithium-ion battery degradation stable modeling and prognosis. *Nat Commun* 2024;15:4332.
- [25] Bischof R, Kraus MA. Multi-objective loss balancing for physics-informed deep learning. *Comput Methods Appl Mech Eng* 2025;439:117914.
- [26] Bizzi A, Nissenbaum L, Pereira JM. Neural conjugate flows: a physics-informed architecture with flow structure. In: Proceedings of the AAAI conference on artificial intelligence; vol. 39. 2025, p. 15576–86.
- [27] Varenberg M. Adjusting for running-in: extension of the archard wear equation. *Tribol Lett* 2022;70(2):59.
- [28] Marian M, Tremmel S. Physics-informed machine learning—an emerging trend in tribology. *Lubricants* 2023;11(11):463.
- [29] Dourado A, Viana F AC. Physics-informed neural networks for corrosion-fatigue prognosis. In: Proceedings of the annual conference of the PHM society; vol. 11. 2019.
- [30] Uzun I, Tolun MR, Sari F, Alpaslan FN. Damage detection in aircraft engine borescope inspection using deep learning. *Neural Comput Appl* 2025;37(24):20105–26.
- [31] Cha Y-J, Choi W, Büyükköztürk O. Deep learning-based crack damage detection using convolutional neural networks. *Comput-Aided Civ Infrastruct Eng* 2017;32(5):361–78. <https://doi.org/10.1111/micc.12263>
- [32] Li L, Liu J, Xing J, Liu Z, Lin K, Du B. Road pothole detection based on crowdsourced data and extended mask r-cnn. *IEEE Trans Intell Transp Syst* 2024.
- [33] Mishra M, Barman T, Ramana GV. Artificial intelligence-based visual inspection system for structural health monitoring of cultural heritage. *J Civil Struct Health Monit* 2024;14(1):103–20. <https://doi.org/10.1007/s13349-022-00643-8>
- [34] Banerjee C, Nguyen K, Fookes C, George K. Physics-informed computer vision: a review and perspectives. *ACM Comput Surv* 2024;57(1). <https://doi.org/10.1145/3689037>
- [35] Huo D, Wang J, Qian Y, Yang Y-H. Glass segmentation with RGB-thermal image pairs. *IEEE Trans Image Process* 2023;32:1911–26.
- [36] Elsayed G, Mahendran A, van Steenkiste S, Greff K, Mozer MC, Kipf T. SAVI + + : Towards end-to-end object-centric learning from real-world videos. *Adv Neural Inf Process Syst* 2022;35:28940–54.
- [37] Gao H, Sun L, Wang J-X. Super-resolution and denoising of fluid flow using physics-informed convolutional neural networks without high-resolution labels. *Phys Fluids* 2021;33(7).
- [38] Cai S, Wang Z, Fuest F, Jeon YJ, Gray C, Karniadakis GE. Flow over an espresso cup: inferring 3-d velocity and pressure fields from tomographic background oriented schlieren via physics-informed neural networks. *J Fluid Mech* 2021;915:A102.
- [39] Cohen T, Welling M. Group equivariant convolutional networks. In: International conference on machine learning. PMLR; 2016, p. 2990–9.
- [40] Gibert X, Patel VM, Chellappa R. Deep multitask learning for railway track inspection. *IEEE Trans Intell Transp Syst* 2017;18(1):153–64. Conference Name: IEEE Transactions on Intelligent Transportation Systems. <https://doi.org/10.1109/TITS.2016.2568758>
- [41] Yucesan YA, Viana F AC. Hybrid physics-informed neural networks for main bearing fatigue prognosis with visual grease inspection. *Comput Ind* 2021;125:103386. <https://doi.org/10.1016/j.compind.2020.103386>
- [42] Yucesan YA, Viana F AC. A hybrid physics-informed neural network for main bearing fatigue prognosis under grease quality variation. *Mech Syst Signal Process* 2022;171:108875. <https://doi.org/10.1016/j.ymsp.2022.108875>
- [43] Zhao X, Gong Z, Zhang Y, Yao W, Chen X. Physics-informed convolutional neural networks for temperature field prediction of heat source layout without labeled data. *Eng Appl Artif Intell* 2023;117:105516. <https://doi.org/10.1016/j.engappai.2022.105516>
- [44] Ji T, Sivakumar AN, Chowdhary G, Driggs-Campbell K. Proactive anomaly detection for robot navigation with multi-sensor fusion. *IEEE Rob Autom Lett* 2022;7(2):4975–82. <https://doi.org/10.1109/LRA.2022.3153989>
- [45] Wang Y, Peng J, Zhang J, Yi R, Wang Y, Wang C. Multimodal industrial anomaly detection via hybrid fusion. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR). 2023, p. 8032–41.
- [46] Nemani V. Hiconstruct\_optimize: Optimized health indicator construction tools. [https://github.com/VNemani14/HIconstruct\\_optimize](https://github.com/VNemani14/HIconstruct_optimize); Accessed: 2025-02-XX.
- [47] Bajarunas K. UnsupervisedHI: Health index estimation using general degradation knowledge. <https://github.com/KBaja/UnsupervisedHI>; Accessed: 2025-02-XX.
- [48] Wang F. Pinn4soh: Physics-informed neural network for lithium-ion battery state-of-health estimation. <https://github.com/wang-fujin/PINN4SOH>; Accessed: 2025-02-XX.
- [49] Wang B, Lei Y, Li N, Li N. A hybrid prognostics approach for estimating remaining useful life of rolling element bearings. *IEEE Trans Reliab* 2018;69(1):401–12.
- [50] Wang F, Zhai Z, Liu B, Di Y, Zhao Y, Tang T. Open access dataset, code library and benchmarking deep learning approaches for state-of-health estimation of lithium-ion batteries. *J Energy Stor* 2024;77:109884. <https://doi.org/10.1016/j.est.2024.109884>
- [51] Saxena A, Goebel K, Simon D, Eklund N. Damage propagation modeling for aircraft engine run-to-failure simulation. In: 2008 International conference on prognostics and health management. 2008, p. 1–9. <https://doi.org/10.1109/PHM.2008.4711414>

- [52] Arias Chao M, Kulkarni C, Goebel K, Fink O. Aircraft engine run-to-failure dataset under real flight conditions for prognostics and diagnostics. *Data* 2021;6(1):5.
- [53] Cubillo A, Perinpanayagam S, Esperon-Míguez M. A review of physics-based models in prognostics: application to gears and bearings of rotating machinery. *Adv Mech Eng* 2016;8(8):1687814016664660. <https://doi.org/10.1177/1687814016664660>
- [54] Reniers JM, Mulder G, Howey DA. Review and performance comparison of mechanical-chemical degradation models for lithium-ion batteries. *J Electrochem Soc* 2019;166(14):A3189–A3200.
- [55] Yu Z-Y, Zhu S-P, Liu Q, Liu Y. A new energy-critical plane damage parameter for multiaxial fatigue life prediction of turbine blades. *Materials* 2017;10(5):513.
- [56] Peng H, Song N, Li F, Tang S. A mechanistic-based data-driven approach for general friction modeling in complex mechanical system. *J Appl Mech* 2022;89(7):071005.
- [57] Liu K, Gebraeel NZ, Shi J. A data-fusion model for developing composite health indices for degradation modeling and prognostic analysis. *IEEE Trans Autom Sci Eng* 2013;10(3):652–64.
- [58] Xu Z, Saleh JH. Machine learning for reliability engineering and safety applications: review of current status and future opportunities. *Reliab Eng Syst Saf* 2021;211:107530.
- [59] Roman D, Saxena S, Robu V, Pecht M, Flynn D. Machine learning pipeline for battery state-of-health estimation. *Nature Mach Intell* 2021;3(5):447–56.
- [60] Wen P, Zhao S, Chen S, Li Y. A generalized remaining useful life prediction method for complex systems based on composite health indicator. *Reliab Eng Syst Saf* 2021;205:107241.
- [61] Frusque G, Nejjar I, Nabavi M, Fink O. Semisupervised health index monitoring with feature generation and fusion. *IEEE Trans Reliab* 2024.
- [62] Liu Y, Hu X, Zhang W. Remaining useful life prediction based on health index similarity. *Reliab Eng Syst Saf* 2019;185:502–10.
- [63] Si X-S. An adaptive prognostic approach via nonlinear degradation modeling: application to battery data. *IEEE Trans Ind Electron* 2015;62(8):5082–96. <https://doi.org/10.1109/TIE.2015.2393840>
- [64] Song C, Liu K. Statistical degradation modeling and prognostics of multiple sensor signals via data fusion: a composite health index approach. *IIESE Trans* 2018;50(10):853–67.
- [65] Liu K, Huang S. Integration of data fusion methodology and degradation modeling process to improve prognostics. *IEEE Trans Autom Sci Eng* 2014;13(1):344–54.
- [66] Bajarunas K, Baptista ML, Goebel K, Chao MA. Health index estimation through integration of general knowledge with unsupervised learning. *Reliab Eng Syst Saf* 2024;251:110352.
- [67] Li Z, Wu J, Yue X. A shape-constrained neural data fusion network for health index construction and residual life prediction. *IEEE Trans Neural Netw Learn Syst* 2020;32(11):5022–33.
- [68] Chen Z, Xia T, Zhou D, Pan E. A health index construction framework for prognostics based on feature fusion and constrained optimization. *IEEE Trans Instrum Meas* 2021;70:1–15.
- [69] Yan T, Wang D, Xia T, Xi L. A generic framework for degradation modeling based on fusion of spectrum amplitudes. *IEEE Trans Autom Sci Eng* 2020;19(1):308–19.
- [70] Kim S, Choi J-H, Kim NH. Data-driven prognostics with low-fidelity physical information for digital twin: physics-informed neural network. *Struct Multidiscip Optim* 2022;65(9):255.
- [71] Chen X, Ma M, Zhao Z, Zhai Z, Mao Z. Physics-informed deep neural network for bearing prognosis with multisensory signals. *J Dyn Monit Diagnost* 2022;200–7.
- [72] Wang F, Du J, Zhao Y, Tang T, Shi J. A deep learning based data fusion method for degradation modeling and prognostics. *IEEE Trans Reliab* 2020;70(2):775–89.
- [73] Qin Y, Yang J, Zhou J, Pu H, Mao Y. A new supervised multi-head self-attention autoencoder for health indicator construction and similarity-based machinery RUL prediction. *Adv Eng Inf* 2023;56:101973. <https://doi.org/10.1016/j.aei.2023.101973>
- [74] Sun T, Yin C, Zheng H, Dong Y. An unsupervised framework for dynamic health indicator construction and its application in rolling bearing prognostics. *Reliab Eng Syst Saf* 2025;260:111039.
- [75] Khaleghi S, Firouz Y, Van Mierlo J, Van Den Bossche P. Developing a real-time data-driven battery health diagnosis method, using time and frequency domain condition indicators. *Appl Energy* 2019;255:113813.
- [76] Caballé NC, Castro IT, Pérez CJ, Lanza-Gutiérrez JM. A condition-based maintenance of a dependent degradation-threshold-shock model in a system with multiple degradation processes. *Reliab Eng Syst Saf* 2015;134:98–109.
- [77] Tajiani B, Vatn J, Naseri M. Optimizing the maintenance threshold in presence of shocks: a numerical framework for systems with non-monotonic degradation. *Reliab Eng Syst Saf* 2024;245:110039.
- [78] Liang Q, Yang Y, Peng C. A reliability model for systems subject to mutually dependent degradation processes and random shocks under dynamic environments. *Reliab Eng Syst Saf* 2023;234:109165.
- [79] Ma T, Weng J, Tian C, Wu M. Change point detection of health stage and remaining useful lifetime prediction for electro-mechanical brake unit on trains. *Measurement* 2024;237:115159.
- [80] Li X, Teng W, Wang L, Hu J, Su Y, Peng D, et al. Trend-constrained pairing based incremental transfer learning for remaining useful life prediction of bearings in wind turbines. *Expert Syst Appl* 2025;263:125731.
- [81] Dong W, Liu S, Bae SJ, Cao Y. Reliability modelling for multi-component systems subject to stochastic deterioration and generalized cumulative shock damages. *Reliab Eng Syst Saf* 2021;205:107260.
- [82] De Felice G, Cini A, Zambon D, Gusev V, Alippi C. Graph-based virtual sensing from sparse and partial multivariate observations. In: *The twelfth international conference on learning representations*. 2024.
- [83] Faghhi Niresi K, Bissig H, Baumann H, Fink O. Physics-enhanced graph neural networks for soft sensing in industrial internet of things. *IEEE Internet Things J* 2024;11(21):34978–90.
- [84] Zhao M, Taal C, Baggerohr S, Fink O. Virtual sensor for real-time bearing load prediction using heterogeneous temporal graph neural networks. In: *PHM society european conference*; vol. 8. 2024, p. 8. <https://doi.org/10.36001/phme.2024.v8i1.3998>
- [85] Niresi KF, Kuhn L, Frusque G, Fink O. Informed graph learning by domain knowledge injection and smooth graph signal representation. In: *2024 32nd european signal processing conference (EUSIPCO)*. IEEE; 2024, p. 2467–71.
- [86] Perera YS, Ratnaweera D, Dasanayaka CH, Abeykoon C. The role of artificial intelligence-driven soft sensors in advanced sustainable process industries: a critical review. *Eng Appl Artif Intell* 2023;121:105988.
- [87] Jiang Y, Yin S, Dong J, Kaynak O. A review on soft sensors for monitoring, control, and optimization of industrial processes. *IEEE Sens J* 2020;21(11):12868–81.
- [88] Kadlec P, Gabrys B, Strandt S. Data-driven soft sensors in the process industry. *Computers & chemical engineering* 2009;33(4):795–814.
- [89] Torgashov A, Skogestad S. The use of first principles model for evaluation of adaptive soft sensor for multicomponent distillation unit. *Chem Eng Res Des* 2019;151:70–8.
- [90] Bikhmukhametov T, Jäschke J. First principles and machine learning virtual flow metering: a literature review. *Journal of Petroleum Science and Engineering* 2020;184:106487.
- [91] Khodarahmi M, Maihami V. A review on kalman filter models. *Arch Comput Methods Eng* 2023;30(1):727–47.
- [92] Alexander R, Campani G, Dinh S, Lima FV. Challenges and opportunities on nonlinear state estimation of chemical and biochemical processes. *Processes* 2020;8(11):1462.
- [93] Cao L, Wang J, Su J, Luo Y, Cao Y, Braatz RD, et al. Comprehensive analysis on machine learning approaches for interpretable and stable soft sensors. *IEEE Trans Instrum Meas* 2025;74:1–17
- [94] Poggi T, Rubagotti M, Bemporad A, Storace M. High-speed piecewise affine virtual sensors. *IEEE Trans Ind Electron* 2011;59(2):1228–37.
- [95] Sefati S, Hegeman R, Alambeigi F, Iordachita I, Armand M. Fbg-based position estimation of highly deformable continuum manipulators: model-dependent vs. data-driven approaches. In: *2019 International symposium on medical robotics (ISMR)*. IEEE; 2019, p. 1–6.
- [96] Wang ZX, He QP, Wang J. Comparison of variable selection methods for PLS-based soft sensor modeling. *J Process Control* 2015;26:56–72.
- [97] Kaneko H, Funatsu K. Application of online support vector regression for soft sensors. *AIChE J* 2014;60(2):600–12.
- [98] Shokri S, Sadeghi MT, Marvast MA, Narasimhan S. Improvement of the prediction performance of a soft sensor model based on support vector regression for production of ultra-low sulfur diesel. *Pet Sci* 2015;12(1):177–88.
- [99] Liu H, Yang C, Huang M, Yoo C. Soft sensor modeling of industrial process data using kernel latent variables-based relevance vector machine. *Appl Soft Comput* 2020;90:106149.
- [100] Liu Y, Chen T, Chen J. Auto-switch Gaussian process regression-based probabilistic soft sensors for industrial multigrade processes with transitions. *Industr Eng Chem Res* 2015;54(18):5037–47.
- [101] Sheng X, Ma J, Xiong W. Smart soft sensor design with hierarchical sampling strategy of ensemble gaussian process regression for fermentation processes. *Sensors* 2020;20(7):1957.
- [102] Sun Q, Ge Z. A survey on deep learning for data-driven soft sensors. *IEEE Trans Ind Inf* 2021;17(9):5853–66.
- [103] Gallareta JG, González-Menorca C, Muñoz P, Vasic MV. Advancements in soft sensor technologies for quality control in process manufacturing: a review. *IEEE Sens J* 2025.
- [104] Wang K, Shang C, Liu L, Jiang Y, Huang D, Yang F. Dynamic soft sensor development based on convolutional neural networks. *Industr Eng Chem Res* 2019;58(26):11521–31.
- [105] Song B, Zhou Y, Shi H, Tao Y, Tan S. A soft sensor for multirate quality variables based on MC-CNN. *IEEE Trans Neural Netw Learn Syst* 2024.
- [106] Ma L, Wang M, Peng K. A novel bidirectional gated recurrent unit-based soft sensor modeling framework for quality prediction in manufacturing processes. *IEEE Sens J* 2022;22(19):18610–19. <https://doi.org/10.1109/JSEN.2022.3199474>
- [107] Tian Y, Xu Y, Zhu Q-X, He Y-L. Novel stacked input-enhanced supervised autoencoder integrated with gated recurrent unit for soft sensing. *IEEE Trans Instrum Meas* 2022;71:1–9.
- [108] Zhang C, Yella J, Huang Y, Qian X, Petrov S, Rzhetsky A, et al. Soft sensing transformer: hundreds of sensors are worth a single word. In: *2021 IEEE International conference on big data (big data)*. IEEE; 2021, p. 1999–2008. <https://doi.org/10.1109/bigdata52589.2021.9671925>
- [109] Yang C, Wang C, Zhao F, Wu B, Fan J-s, Zhang Y. Smart virtual sensing for deep excavations using real-time ensemble graph neural networks. *Autom Constr* 2025;172:106040.
- [110] Wang J, Sui E, Wang W, Zhou X, Zhang Z, Li J. A stable soft sensor based on causal inference and graph convolutional network for batch processes. *Expert Syst Appl* 2025;263:125692.
- [111] Hu W, Wang Y, He Y, Qian L, Zhang D, Meng Y. Soft sensing technique for mass customization based on heterogeneous causal graph attention networks. *Adv Eng Inf* 2025;65:103139.
- [112] Liu Y, Jia M, Xu D, Yang T, Yao Y. Physics-guided graph learning soft sensor for chemical processes. *Chemometr Intell Laborat Syst* 2024;249:105131.

- [113] Aina W, Pan Q, Yunbo Y, Guang Z, Ximing S. Physically informed hierarchical learning based soft sensing for aero-engine health management unit. *Chin J Aeronaut* 2025;38(3):103308.
- [114] Funatsu K. Soft sensors: chemoinformatic model for efficient control and operation in chemical plants. *Mol Inform* 2016;35(11-12):549–54.
- [115] Kral C, Haumer A, Lee SB. A practical thermal model for the estimation of permanent magnet and stator winding temperatures. *IEEE Trans Power Electron* 2013;29(1):455–64.
- [116] Kim H, Jung H-S, Sul S-K. Stator winding temperature and magnet temperature estimation of IPMSM based on high-frequency voltage signal injection. *IEEE Trans Ind Electron* 2022;70(3):2296–306.
- [117] Dini P, Colicelli A, Saponara S. Review on modeling and soc/soh estimation of batteries for automotive applications. *Batteries* 2024;10(1):34.
- [118] Darvishi H, Ciuonzo D, Eide ER, Rossi PS. Sensor-fault detection, isolation and accommodation for digital twins via modular data-driven architecture. *IEEE Sens J* 2020;21(4):4827–38.
- [119] Frafjord AJ, Radicke J-P, Keprate A, Komulainen TM. Data-driven approaches for deriving a soft sensor in a district heating network. *Energy* 2024;292:130426.
- [120] Nie L, Ren Y, Wu R, Tan M. Sensor fault diagnosis, isolation, and accommodation for heating, ventilating, and air conditioning systems based on soft sensor. In: *Actuators*; vol. 12. MDPI; 2023, p. 389.
- [121] Fey M, Lenssen JE. Fast graph representation learning with pytorch geometric. In: *ICLR Workshop on representation learning on graphs and manifolds*. 2019.
- [122] Fey M, Sunil J, Nitta A, Puri R, Shah M, Stojanovic B, et al. PyG 2.0: Scalable learning on real world graphs. In: *Temporal graph learning workshop @ KDD*. 2025.
- [123] Cini A, Marisca I. Torch spatiotemporal 2022; <https://github.com/TorchSpatiotemporal/tsl>.
- [124] Tong S, Guo R, Wang W, Tian X, Wei L, Zhang L, et al. A soft sensor method with uncertainty-awareness and self-explanation based on large language models enhanced by domain knowledge retrieval. 2025, arXiv preprint arXiv:250103295.
- [125] Niresi KF, Nejjar I, Fink O. Efficient unsupervised domain adaptation regression for spatial-temporal sensor fusion. *IEEE Internet Things J* 2025.
- [126] Orsagh RF, Sheldon J, Klenke CJ. Prognostics/diagnostics for gas turbine engine bearings; vol. 36843. 2003.
- [127] Zong H, Wang H, Hong R. Dynamic research on single-row slewing bearing with local spalling of inner ring. *Adv Mechan Eng* 2016;8(2):1687814016632488.
- [128] Tadini M, Boltezar M. Improved model of a ball bearing for the simulation of vibration signals due to faults during run-up. *J Sound Vib* 2011;330(17):4287–301.
- [129] Singh S, Howard CQ, Hansen CH. An extensive review of vibration modelling of rolling element bearings with localised and extended defects. *J Sound Vib* 2015;357:300–30.
- [130] Li CJ, Lee H. Gear fatigue crack prognosis using embedded model, gear dynamic model and fracture mechanics. *Mech Syst Signal Process* 2005;19(4):836–46.
- [131] Cubillo A, Perinpanayagam S, Esperon-Miguez M. A review of physics-based models in prognostics: application to gears and bearings of rotating machinery. *Adv Mechan Eng* 2016;8(8):1687814016664660.
- [132] Daigle MJ, Goebel K. A model-based prognostics approach applied to pneumatic valves. *Int J Progn Health Manage Volume 2 (color)* 2011;84.
- [133] Daigle MJ, Goebel K. Model-based prognostics with concurrent damage progression processes. *IEEE Trans Syst Man Cybernet Syst* 2012;43(3):535–46.
- [134] Ma S, Flanigan KA, Bergés M. Bridging the reality gap in digital twins with context-aware, physics-guided deep learning. 2025. 2505.11847.
- [135] Lei Y, Li N, Guo L, Li N, Yan T, Lin J. Machinery health prognostics: a systematic review from data acquisition to RUL prediction. *Mech Syst Signal Process* 2018;104:799–834.
- [136] Khan S, Yairi T. A review on the application of deep learning in system health management. *Mech Syst Signal Process* 2018;107:241–65.
- [137] Hagemeyer S, Zeiler P, Huber MF. On the integration of fundamental knowledge about degradation processes into data-driven diagnostics and prognostics using theory-guided data science. In: *PHM society european conference*; vol. 7. 2022, p. 156–65.
- [138] Eklund N HW. Using synthetic data to train an accurate real-world fault detection system. In: *The proceedings of the multiconference on "computational engineering in systems applications"*; vol. 1. IEEE; 2006, p. 483–8.
- [139] Sobie C, Freitas C, Nicolai M. Simulation-driven machine learning: bearing fault classification. *Mech Syst Signal Process* 2018;99:403–19.
- [140] Thelen A, Zhang X, Fink O, Lu Y, Ghosh S, Youn BD, et al. A comprehensive review of digital twin-part 1: modeling and twinning enabling technologies. *Struct Multidiscip Optim* 2022;65(12):354.
- [141] Gao Y, Liu X, Xiang J. Fem simulation-based generative adversarial networks to detect bearing faults. *IEEE Trans Ind Inf* 2020;16(7):4961–71.
- [142] D'Urso D, Chiacchio F, Cavalieri S, Gambadoro S, Khodayee SM. Predictive maintenance of standalone steel industrial components powered by a dynamic reliability digital twin model with artificial intelligence. *Reliab Eng Syst Saf* 2024;243:109859.
- [143] Teubert C, Griffith KJ, Corbetta M, Kulkarni C, Banerjee P, Watkins J, et al. ProgPy Python Prognostics Packages. 2025. <https://doi.org/10.5281/ZENODO.8097013>
- [144] Jiang J, Chen B, Fu B, Long M. Transfer-learning-library. <https://github.com/thuml/Transfer-Learning-Library>; 2020.
- [145] Daigle M, Kulkarni CS. Electrochemistry-based battery modeling for prognostics. In: *Annual conference of the PHM society*; vol. 5. 2013.
- [146] Bole B, Kulkarni CS, Daigle M. Adaptation of an electrochemistry-based li-ion battery model to account for deterioration observed under randomized use. In: *Annual conference of the PHM society*; vol. 6. 2014.
- [147] Frank S, Heaney M, Jin X, Robertson J, Cheung H, Elmore R, et al. Hybrid model-based and data-driven fault detection and diagnostics for commercial buildings. *Tech. Rep.*; National Renewable Energy Lab.(NREL), Golden, CO (United States); 2016.
- [148] Huber LG, Palmé T, Chao MA. Physics-informed machine learning for predictive maintenance: applied use-cases. In: *2023 10th IEEE swiss conference on data science (SDS)*. IEEE; 2023, p. 66–72.
- [149] Randall RB, Antoni J, Chobsaard S. The relationship between spectral correlation and envelope analysis in the diagnostics of bearing faults and other cyclostationary machine signals. *Mech Syst Signal Process* 2001;15(5):945–62.
- [150] Biggio L, Bendinelli T, Kulkarni C, Fink O. Ageing-aware battery discharge prediction with deep learning. *Appl Energy* 2023;346:121229.
- [151] Chakraborty S. Transfer learning based multi-fidelity physics informed deep neural network. *J Comput Phys* 2021;426:109942.
- [152] Schröder L, Dimitrov NK, Verelst DR, Sørensen JA. Using transfer learning to build physics-informed machine learning models for improved wind farm monitoring. *Physics* 2022;15(2):558.
- [153] Liu C, Ricardo Mauricio A, Qi J, Peng D, Gryllias K. Domain adaptation digital twin for rolling element bearing prognostics. *Online Proceedings of PHM2020* 2020;1–10.
- [154] Wang Q, Taal C, Fink O. Integrating expert knowledge with domain adaptation for unsupervised fault diagnosis. *IEEE Trans Instrum Meas* 2021;71:1–12.
- [155] Geegel O, Dias JP, Ekwaro-Osire S, Alves DS, Machado TH, Daniel GB, et al. Simulation-driven deep learning approach for wear diagnostics in hydrodynamic journal bearings. *J Tribol* 2021;143(8):084501.
- [156] Khan A, Hwang H, Kim HS. Synthetic data augmentation and deep learning for the fault diagnosis of rotating machines. *Mathematics* 2021;9(18):2336.
- [157] Jan Z, Ahamed F, Mayer W, Patel N, Grossmann G, Stumptner M, et al. Artificial intelligence for industry 4.0: systematic review of applications, challenges, and opportunities. *Expert Syst Appl* 2023;216:119456.
- [158] Lu Y, Wang B, Zhao Y, Yang X, Li L, Dong M, et al. Physics-informed surrogate modeling for hydro-fracture geometry prediction based on deep learning. *Energy* 2022;253:124139.
- [159] Suttakul P, Vo D, Fongsamoot T, Wanison R, Mona Y, Katongtung T, et al. The role of machine learning for insight into the material behavior of lattices: a surrogate model based on data from finite element simulation. *Res Eng* 2024;23:102547.
- [160] Torzoni M, Manzoni A, Mariani S. A multi-fidelity surrogate model for structural health monitoring exploiting model order reduction and artificial neural networks. *Mech Syst Signal Process* 2023;197:110376.
- [161] Bommasani R, Hudson DA, Adeli E, Altman R, Arora S, von Arx S, et al. On the opportunities and risks of foundation models. 2022. 2108.07258.
- [162] Zheng S, Pan K, Liu J, Chen Y. Empirical study on fine-tuning pre-trained large language models for fault diagnosis of complex systems. *Reliab Eng Syst Saf* 2024;252:110382.
- [163] Abramson J, Adler J, Dunger J, Evans R, Green T, Pritzel A, et al. Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature* 2024;630(8016):493–500.
- [164] Bodnar C, Bruinisma WP, Lucic A, Stanley M, Brandstetter J, Garvan P, et al. Aurora: a foundation model of the atmosphere. *CoRR* 2024.
- [165] Price I, Sanchez-Gonzalez A, Alet F, Andersson TR, El-Kadi A, Masters D, et al. Probabilistic weather forecasting with machine learning. *Nature* 2025;637(8044):84–90.
- [166] Subramanian S, Harrington P, Keutzer K, Bhimji W, Morozov D, Mahoney MW, et al. Towards foundation models for scientific machine learning: characterizing scaling and transfer behavior. *Adv Neural Inf Process Syst* 2023;36:71242–62.
- [167] Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, et al. Language models are few-shot learners. *Adv Neural Inf Process Syst* 2020;33:1877–901.
- [168] Herde M, Raoniç B, Rohner T, Käppeli R, Molinaro R, de Bézenac E, et al. Poseidon: Efficient foundation models for PDEs. 2024. 2405.19101.
- [169] Sharma V, O'donn RT, Tesini P, Ravesloot J, Taal C, Fink O. Equi-euler graphnet: an equivariant, temporal-dynamics informed graph neural network for dual force and trajectory prediction in multi-body systems. *Mech Syst Signal Process* 2025;241:113533. <https://doi.org/10.1016/j.ymsp.2025.113533>
- [170] Akyürek E, Schuurmans D, Andreas J, Ma T, Zhou D. What learning algorithm is in-context learning? investigations with linear models. In: *The eleventh international conference on learning representations*. 2023.
- [171] Bai Y, Chen F, Wang H, Xiong C, Mei S. Transformers as statisticians: provable in-context learning with in-context algorithm selection. *Adv Neural Inf Process Syst* 2023;36:57125–211.
- [172] Nejjar I, Ahmed F, Fink O. IM-Context: in-context learning for imbalanced regression tasks. *Trans Mach Learn Res* 2024;.
- [173] Ruthotto L, Haber E. An introduction to deep generative modeling. *GAMM-Mitteilungen* 2021;44(2):e202100008.
- [174] Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, et al. Generative adversarial nets. *Adv Neural Inf Process Syst* 2014;27.
- [175] Hinton GE, Krizhevsky A, Wang SD. Transforming auto-encoders. In: *International conference on artificial neural networks*. Springer; 2011, p. 44–51.
- [176] Kingma DP, Welling M, et al. An introduction to variational autoencoders. *Front Trends Mach Learn* 2019;12(4):307–92.
- [177] Rezende D, Mohamed S. Variational inference with normalizing flows. In: *International conference on machine learning*. PMLR; 2015, p. 1530–8.
- [178] Sohl-Dickstein J, Weiss E, Maheswaranathan N, Ganguli S. Deep unsupervised learning using nonequilibrium thermodynamics. In: *International conference on machine learning*. PMLR; 2015, p. 2256–65.
- [179] Tomczak JM. *Deep Generative Modeling*. Cham: Springer International Publishing; 2024. <https://doi.org/10.1007/978-3-031-64087-2>

- [180] Linder-Norén E. eriklindernoren/PyTorch-GAN. 2018. original-date: 2018-04-21T13:06:40Z.
- [181] Yoon J. TimeGAN: Time-series generative adversarial networks codebase. <https://github.com/jsoo0823/TimeGAN>; 2019.
- [182] Li D. Mad-gan: Multivariate anomaly detection with gans implementation. <https://github.com/LiDan456/MAD-GAN>; 2019.
- [183] Zhao Y. PyOD: A python toolbox for scalable outlier detection. <https://github.com/yzhao062/pyod>; 2019.
- [184] Science A. Time series diffusion (TSDiff) repository. <https://github.com/amazon-science/time-series-diffusion>; 2023.
- [185] Stief A, Tan R, Cao Y, Ottewill JR, Thornhill NF, Baranowski J. A heterogeneous benchmark dataset for data analytics: multiphase flow facility case study. *J Process Control* 2019;79:41–55. <https://doi.org/10.1016/j.jprocont.2019.04.009>
- [186] Nguyen R, Singh SK, Rai R. Physics-infused fuzzy generative adversarial network for robust failure prognosis. *Mech Syst Signal Process* 2023;184:109611.
- [187] Böck B, Oeldemann A, Mayer T, Rossetto F, Utschick W. Physics-informed generative modeling of wireless channels. In: Forty-second international conference on machine learning. 2025.
- [188] Zhou T, Drogue EL, Mosleh A. Physics-informed deep learning: a promising technique for system reliability assessment. *Appl Soft Comput* 2022;126:109217.
- [189] Battaglia P, Pascanu R, Lai M, Jimenez Rezende D, et al. Interaction networks for learning about objects, relations and physics. *Adv Neural Inf Process Syst* 2016;29.
- [190] Greydanus S, Dzamba M, Yosinski J. Hamiltonian neural networks. *Adv Neural Inf Process Syst* 2019;32.
- [191] Bastek J-H, Sun W, Kochmann D. Physics-informed diffusion models. In: The thirteenth international conference on learning representations. 2025.
- [192] Qiu J, Huang J, Zhang X, Lin Z, Pan M, Liu Z, et al. Pi-fusion: physics-informed diffusion model for learning fluid dynamics. *CoRR* 2024;abs/2406.03711.
- [193] Rissanen S, Heinonen M, Solin A. Generative modelling with inverse heat dissipation. In: The eleventh international conference on learning representations. 2023.
- [194] Liu Y, Jiang H, Yao R, Zhu H. Interpretable data-augmented adversarial variational autoencoder with sequential attention for imbalanced fault diagnosis. *J Manuf Syst* 2023;71:342–59.
- [195] Zhao X, Yao J, Deng W, Jia M, Liu Z. Normalized conditional variational autoencoder with adaptive focal loss for imbalanced fault diagnosis of bearing-rotor system. *Mech Syst Signal Process* 2022;170:108826.
- [196] Tian J, Jiang Y, Zhang J, Luo H, Yin S. A novel data augmentation approach to fault diagnosis with class-imbalance problem. *Reliab Eng Syst Saf* 2024;243:109832.
- [197] Costa N, Sanchez L. Few-shot generative compression approach for system health monitoring. *Reliab Eng Syst Saf* 2025;256:110749.
- [198] Hong D, Kim B. 1D convolutional neural network-based adaptive algorithm structure with system fault diagnosis and signal feature extraction for noise and vibration enhancement in mechanical systems. *Mech Syst Signal Process* 2023;197:110395.
- [199] Rombach K, Michau G, Fink O. Controlled generation of unseen faults for partial and open-partial domain adaptation. *Reliab Eng Syst Saf* 2023;230:108857.
- [200] Sun H, Cao Y, Dong H, Fink O. Unseen visual anomaly generation. In: Proceedings of the computer vision and pattern recognition conference. 2025, p. 25508–17.
- [201] Ren H, Dai Z, Zhou B, Zhang B, Yin A, Pei L, et al. Vibration monitoring and semisupervised multivariate invertible deep probabilistic learning for gearbox faults identification. *IEEE Sens J* 2022;22(22):22020–9.
- [202] Luo Q, Chen J, Zi Y, Chang Y, Feng Y. Multi-mode non-Gaussian variational autoencoder network with missing sources for anomaly detection of complex electromechanical equipment. *ISA Trans* 2023;134:144–58.
- [203] Hu W, Frusque G, Wang T, Chu F, Fink O. Classifier-free diffusion-based weakly-supervised approach for health indicator derivation in rotating machines: advancing early fault detection and condition monitoring. *Reliab Eng Syst Saf* 2025;111397.
- [204] Zhu G, Zhao H, Liu H, Sun H. A novel LSTM-GAN algorithm for time series anomaly detection. In: 2019 Prognostics and system health management conference (PHM-Qingdao). IEEE; 2019, p. 1–6.
- [205] Fang J, Yang F, Tong R, Yu Q, Dai X. Fault diagnosis of electric transformers based on infrared image processing and semi-supervised learning. *Global Energy Interconnection* 2021;4(6):596–607.
- [206] Xu Z, Guo Y, Saleh JH. A physics-informed dynamic deep autoencoder for accurate state-of-health prediction of lithium-ion battery. *Neural Comput Appl* 2022;34(18):15997–6017.
- [207] Xiong J, Fink O, Zhou J, Ma Y. Controlled physics-informed data generation for deep learning-based remaining useful life prediction under unseen operation conditions. *Mech Syst Signal Process* 2023;197:110359.
- [208] Gao S, Liu J, Zhang Z, Feng C, He B, Zio E. Physics-guided generative adversarial networks for fault detection of underwater thruster. *Ocean Eng* 2023;286:115585.
- [209] Sanchez L, Costa N, Couso I. Addressing data scarcity in industrial reliability assessment with physically informed echo state networks. *Reliab Eng Syst Saf* 2025;111135.
- [210] Hollmann N, Müller S, Purucker L, Krishnakumar A, Körfer M, Hoo SB, et al. Accurate predictions on small data with a tabular foundation model. *Nature* 2025;637(8045):319–26. <https://doi.org/10.1038/s41586-024-08328-6>
- [211] Hennebold C, Krauß J, Huber MF. Synthetic data generation using causal models for injection molding processes. *Procedia CIRP* 2025;134:831–6. <https://doi.org/10.1016/j.procir.2025.02.202>
- [212] Wu D, Qiu F, Xie Y. Doflow: flow-based generative models for interventional and counterfactual forecasting on time series. In: Submitted to the fourteenth international conference on learning representations. 2025, under review.
- [213] Cheng Y, Wang Z, Xiao T, Zhong Q, Suo J, He K. CausalTime: realistically generated time-series for benchmarking of causal discovery. 2023.
- [214] Gamella JL, Peters J, Bühlmann P. Causal chambers as a real-world physical testbed for AI methodology. *Nature Mach Intell* 2025;7(1):107–18. Publisher: Nature Publishing Group. <https://doi.org/10.1038/s42256-024-00964-x>
- [215] Khaleghi B, Khamis A, Karray FO, Razavi SN. Multisensor data fusion: a review of the state-of-the-art. *Inf fusion* 2013;14(1):28–44.
- [216] Meng T, Jing X, Yan Z, Pedrycz W. A survey on machine learning for data fusion. *Inf Fusion* 2020;57:115–29. <https://doi.org/10.1016/j.inffus.2019.12.001>
- [217] Jardine A KS, Lin D, Banjevic D. A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mech Syst Signal Process* 2006;20(7):1483–510.
- [218] Jose S, Nguyen K TP, Medjaher K. Multimodal Machine Learning in Prognostics and Health Management of Manufacturing Systems. Cham: Springer International Publishing. ISBN 978-3-031-30510-8; 2023, p. 167–97. [https://doi.org/10.1007/978-3-031-30510-8\\_9](https://doi.org/10.1007/978-3-031-30510-8_9)
- [219] De Jong Y, Velasco-Hernandez G, Barry J, Walsh J. Sensor and sensor fusion technology in autonomous vehicles: a review. *Sensors* 2021;21(6):2140.
- [220] Dong H, Gu W, Zhang X, Xu J, Ai R, Lu H, et al. Superfusion: multilevel lidar-camera fusion for long-range hd map generation. In: 2024 IEEE International conference on robotics and automation (ICRA). IEEE; 2024, p. 9056–62.
- [221] Chen X, Zhang T, Wang Y, Wang Y, Zhao H. Futr3d: a unified sensor fusion framework for 3d detection. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 2023, p. 172–81.
- [222] Dong H, Chen X, Särkkä S, Stachniss C. Online pole segmentation on range images for long-term lidar localization in urban environments. *Rob Auton Syst* 2023;159:104283.
- [223] King RC, Villeneuve E, White RJ, Sherratt RS, Holderbaum W, Harwin WS. Application of data fusion techniques and technologies for wearable health monitoring. *Medical engineering & physics* 2017;42:1–12.
- [224] Xu J, Wang Y, Xu L. Phm-oriented integrated fusion prognostics for aircraft engines based on sensor data. *IEEE Sens J* 2014;14(4):1124–32. <https://doi.org/10.1109/JSEN.2013.2293517>
- [225] Sarkar S, Sarkar S, Mukherjee K, Ray A, Srivastav A. Multi-sensor information fusion for fault detection in aircraft gas turbine engines. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering* 2013;227(12):1988–2001.
- [226] Xia J, Huang R, Chen Z, He G, Li W. A novel digital twin-driven approach based on physical-virtual data fusion for gearbox fault diagnosis. *Reliab Eng Syst Saf* 2023;240:109542.
- [227] El Hadraoui H, Ouahabi N, El Bazi N, Laayati O, Zegrari M, Chebak A. Toward an intelligent diagnosis and prognostic health management system for autonomous electric vehicle powertrains: a novel distributed intelligent digital twin-based architecture. *IEEE Access* 2024.
- [228] Liang PP, Zadeh A, Morency L-P. Foundations & trends in multimodal machine learning: principles, challenges, and open questions. *ACM Comput Surv* 2024;56(10):1–42.
- [229] Bonetto E, Xu C, Ahmad A. Grade: generating realistic and dynamic environments for robotics research with isaac sim. *Int J Rob Res* 2025;02783649251346211. <https://doi.org/10.1177/02783649251346211>
- [230] Sun D, Li Y, Liu Z, Jia S, Noman K. Physics-inspired multimodal machine learning for adaptive correlation fusion based rotating machinery fault diagnosis. *Inf Fusion* 2024;108:102394. <https://doi.org/10.1016/j.inffus.2024.102394>
- [231] Wang T, Gan V JL. Automated joint 3d reconstruction and visual inspection for buildings using computer vision and transfer learning. *Autom Constr* 2023;149:104810. <https://doi.org/10.1016/j.autcon.2023.104810>
- [232] Xu C, Mielle M, Laborde A, Waseem A, Forest F, Fink O. Exploiting semantic scene reconstruction for estimating building envelope characteristics. *Build Environ* 2025;275:112731. <https://doi.org/10.1016/j.buildenv.2025.112731>
- [233] Hu P, Zhong Q, Zheng S, Chen X, Peng L. 3D reconstruction and depth estimation method for local anomalies of rail surface based on multi-view stereo matching. *Meas Sci Technol* 2024;36(1):015004.
- [234] Bhavanasi G, Neven D, Artega M, Ditzel S, Dehaeck S, Bey-Temsamani A. Enhanced vision-based quality inspection: a multiview artificial intelligence framework for defect detection. *Sensors* 2025;25(6):1703.
- [235] Cao X, Li Y, Yu K, Zhang Y. Integrated multimodal fault diagnosis of industrial gearboxes using vibration signals and infrared thermal imaging. In: 2024 Global reliability and prognostics and health management conference (PHM-Beijing). 2024, p. 1–5. <https://doi.org/10.1109/PHM-Beijing63284.2024.10874640>
- [236] Yin S, Fu C, Zhao S, Li K, Sun X, Xu T, et al. A survey on multimodal large language models. *Natl Sci Rev* 2024;11(12).
- [237] Jose S, Nguyen K TP, Medjaher K, Zemouri R, Lévesque M, Tahan A. Advancing multimodal diagnostics: integrating industrial textual data and domain knowledge with large language models. *Expert Syst Appl* 2024;255:124603. <https://doi.org/10.1016/j.eswa.2024.124603>
- [238] Luo RC, Yih C-C, Su KL. Multisensor fusion and integration: approaches, applications, and future research directions. *IEEE Sens J* 2002;2(2):107–19.
- [239] Castanedo F. A review of data fusion techniques. *The Scientific World Journal* 2013;2013(1):704504. <https://doi.org/10.1155/2013/704504>
- [240] Zhao F, Zhang C, Geng B. Deep multimodal data fusion. *ACM Comput Surv* 2024;56(9). <https://doi.org/10.1145/3649447>
- [241] Li S, Kang X, Fang L, Hu J, Yin H. Pixel-level image fusion: a survey of the state of the art. *Inf Fusion* 2017;33:100–12. <https://doi.org/10.1016/j.inffus.2016.05.004>
- [242] Pajares G, De La Cruz JM. A wavelet-based image fusion tutorial. *Pattern Recognit* 2004;37(9):1855–72.
- [243] Rubinstein R, Bruckstein AM, Elad M. Dictionaries for sparse representation modeling. *Proc IEEE* 2010;98(6):1045–57.

- [244] Urrea C, Agramonte R. Kalman filter: historical overview and review of its use in robotics 60 years after its creation. *J Sensors* 2021;2021(1):9674015.
- [245] Hussain M, O'Nils M, Lundgren J, Mousavirad SJ. A comprehensive review on deep learning-based data fusion. *IEEE Access* 2024;12:180093–124. <https://doi.org/10.1109/ACCESS.2024.3508271>
- [246] Baltrušaitis T, Ahuja C, Morency L-P. Multimodal machine learning: a survey and taxonomy. *IEEE Trans Pattern Anal Mach Intell* 2019;41(2):423–43. <https://doi.org/10.1109/TPAMI.2018.2798607>
- [247] Ayman A, Onsy A, Attallah O, Brooks H, Morsi I. Feature learning for bearing prognostics: a comprehensive review of machine/deep learning methods, challenges, and opportunities. *Measurement* 2025;245:116589. <https://doi.org/10.1016/j.measurement.2024.116589>
- [248] Gao J, Li P, Chen Z, Zhang J. A survey on deep learning for multimodal data fusion. *Neural Comput* 2020;32(5):829–64. [https://doi.org/10.1162/neco\\_a\\_01273](https://doi.org/10.1162/neco_a_01273)
- [249] Yin Y, Tian J, Liu X. Remaining useful life prediction based on parallel multi-scale feature fusion network. *J Intell Manuf* 2025;36(5):3111–27.
- [250] Dang HV, Tran-Ngoc H, Nguyen TV, Bui-Tien T, De Roeck G, Nguyen HX. Data-driven structural health monitoring using feature fusion and hybrid deep learning. *IEEE Trans Autom Sci Eng* 2021;18(4):2087–103. <https://doi.org/10.1109/TASE.2020.3034401>
- [251] Xu W, Jiang Q, Shen Y, Zhu Q, Xu F. New RUL prediction method for rotating machinery via data feature distribution and spatial attention residual network. *IEEE Trans Instrum Meas* 2023;72:1–9. <https://doi.org/10.1109/TIM.2023.3246526>
- [252] Theiler R, Fink O. Heterogeneous graph neural networks for short-term state forecasting in power systems across domains and time scales: A hydroelectric power plant case study. 2025. 2507.06694.
- [253] Pu H, Chen Z, Liu J, Yang X, Ren C, Liu H, et al. Research on decision-level fusion method based on structural causal model in system-level fault detection and diagnosis. *Eng Appl Artif Intell* 2023;126:107095. <https://doi.org/10.1016/j.engappai.2023.107095>
- [254] Kibrete F, Engida Woldemichael D, Shimels Gebremedhen H. Multi-sensor data fusion in intelligent fault diagnosis of rotating machines: a comprehensive review. *Measurement* 2024;232:114658. <https://doi.org/10.1016/j.measurement.2024.114658>
- [255] Li X, Jiang H, Niu M, Wang R. An enhanced selective ensemble deep learning method for rolling bearing fault diagnosis with beetle antennae search algorithm. *Mech Syst Signal Process* 2020;142:106752. <https://doi.org/10.1016/j.ymsp.2020.106752>
- [256] Ji X, Ren Y, Tang H, Shi C, Xiang J. An intelligent fault diagnosis approach based on Dempster-Shafer theory for hydraulic valves. *Measurement* 2020;165:108129. <https://doi.org/10.1016/j.measurement.2020.108129>
- [257] Xu BG. Intelligent fault inference for rotating flexible rotors using Bayesian belief network. *Expert Syst Appl* 2012;39(1):816–22. <https://doi.org/10.1016/j.eswa.2011.07.079>
- [258] Zhong J-H, Wong PK, Yang Z-X. Fault diagnosis of rotating machinery based on multiple probabilistic classifiers. *Mech Syst Signal Process* 2018;108:99–114. <https://doi.org/10.1016/j.ymsp.2018.02.009>
- [259] Wang Z-Y, Lu C, Zhou B. Fault diagnosis for rotary machinery with selective ensemble neural networks. *Mech Syst Signal Process* 2018;113:112–30. *SI: IMETI-MechElectro*. <https://doi.org/10.1016/j.ymsp.2017.03.051>
- [260] Song D, Ma T, Li Y, Xu F. Data and decision level fusion-based crack detection for compressor blade using acoustic and vibration signal. *IEEE Sens J* 2022;22(12):12209–18. <https://doi.org/10.1109/JSEN.2022.3173924>
- [261] Luwei KC, Yunusa-Kaltungo A, Sha'aban YA. Integrated fault detection framework for classifying rotating machine faults using frequency domain data fusion and artificial neural networks. *Machines* 2018;6(4):59.
- [262] Liu K, Gebrael NZ, Shi J. A data-level fusion model for developing composite health indices for degradation modeling and prognostic analysis. *IEEE Trans Autom Sci Eng* 2013;10(3):652–64. <https://doi.org/10.1109/TASE.2013.2250282>
- [263] Liu K, Huang S. Integration of data fusion methodology and degradation modeling process to improve prognostics. *IEEE Trans Autom Sci Eng* 2016;13(1):344–54. <https://doi.org/10.1109/TASE.2014.2349733>
- [264] Chen S, Wen P, Zhao S, Huang D, Wu M, Zhang Y. A data fusion-based methodology of constructing health indicators for anomaly detection and prognostics. In: 2018 International conference on sensing, diagnostics, prognostics, and control (SDPC). 2018, p. 570–6. <https://doi.org/10.1109/SDPC.2018.8664723>
- [265] Song C, Liu K, Zhang X. Integration of data-level fusion model and kernel methods for degradation modeling and prognostic analysis. *IEEE Trans Reliab* 2018;67(2):640–50. <https://doi.org/10.1109/TR.2017.2715180>
- [266] Wang D, Liu K. An integrated deep learning-based data fusion and degradation modeling method for improving prognostics. *IEEE Trans Autom Sci Eng* 2024;21(2):1713–26. <https://doi.org/10.1109/TASE.2023.3242355>
- [267] Lim P, Goh CK, Tan KC, Dutta P. Multimodal degradation prognostics based on switching Kalman filter ensemble. *IEEE Trans Neural Netw Learn Syst* 2017;28(1):136–48. <https://doi.org/10.1109/TNNLS.2015.2504389>
- [268] Li G, Wei J, He J, Yang H, Meng F. Implicit Kalman filtering method for remaining useful life prediction of rolling bearing with adaptive detection of degradation stage transition point. *Reliab Eng Syst Saf* 2023;235:109269. <https://doi.org/10.1016/j.res.2023.109269>
- [269] Wang Y, Jiang K, Xu L. Sensor-data-driven fusion prognostic framework for complex engineered systems. *IEEE Sens J* 2021;21(18):20421–30. <https://doi.org/10.1109/JSEN.2021.3095556>
- [270] Lee M-S, Shifat TA, Hur J-W. Kalman filter assisted deep feature learning for RUL prediction of hydraulic gear pump. *IEEE Sens J* 2022;22(11):11088–97. <https://doi.org/10.1109/JSEN.2022.3167926>
- [271] Kalman RE. A new approach to linear filtering and prediction problems 1960.
- [272] Gordon NJ, Salmond DJ, Smith A FM. Novel approach to nonlinear/non-gaussian Bayesian state estimation. In: *IEE Proceedings f (radar and signal processing)*; vol. 140. IET; 1993, p. 107–13.
- [273] Peng S, Li X, Zhang S, Jian Z, Duan H, and ZW. A review: state estimation based on hybrid models of Kalman filter and neural network. *Syst Sci Control Eng* 2023;11(1):2173682. <https://doi.org/10.1080/21642583.2023.2173682>
- [274] Gebrael N, Lawley M, Liu R, Parmeshwaran V. Residual life predictions from vibration-based degradation signals: a neural network approach. *IEEE Trans Ind Electron* 2004;51(3):694–700. <https://doi.org/10.1109/TIE.2004.824875>
- [275] Tian Z. An artificial neural network method for remaining useful life prediction of equipment subject to condition monitoring. *J Intell Manuf* 2012;23(2):227–37.
- [276] Sateesh Babu G, Zhao P, Li X-L. Deep convolutional neural network based regression approach for estimation of remaining useful life. In: *International conference on database systems for advanced applications*. Springer; 2016, p. 214–28.
- [277] Li X, Zhang W, Ma H, Luo Z, Li X. Degradation alignment in remaining useful life prediction using deep cycle-consistent learning. *IEEE Trans Neural Netw Learn Syst* 2022;33(10):5480–91. <https://doi.org/10.1109/TNNLS.2021.3070840>
- [278] Zheng S, Ristovski K, Farahat A, Gupta C. Long short-term memory network for remaining useful life estimation. In: 2017 IEEE International conference on prognostics and health management (ICPHM). 2017, p. 88–95. <https://doi.org/10.1109/ICPHM.2017.7998311>
- [279] Ma M, Mao Z. Deep-convolution-based LSTM network for remaining useful life prediction. *IEEE Trans Ind Inf* 2021;17(3):1658–67. <https://doi.org/10.1109/TII.2020.2991796>
- [280] Cordoni F, Bacchiega G, Bondani G, Radu R, Muradore R. A multi-modal unsupervised fault detection system based on power signals and thermal imaging via deep autoencoder neural network. *Eng Appl Artif Intell* 2022;110:104729. <https://doi.org/10.1016/j.engappai.2022.104729>
- [281] Zhang Y, Ding J, Li Y, Ren Z, Feng K. Multi-modal data cross-domain fusion network for gearbox fault diagnosis under variable operating conditions. *Eng Appl Artif Intell* 2024;133:108236. <https://doi.org/10.1016/j.engappai.2024.108236>
- [282] Miao M, Yu J. Deep feature interactive network for machinery fault diagnosis using multi-source heterogeneous data. *Reliab Eng Syst Saf* 2024;242:109795. <https://doi.org/10.1016/j.res.2023.109795>
- [283] Li Y, Yu F, Cai Q, Yuan K, Wan R, Li X, et al. Image fusion of fault detection in power system based on deep learning. *Cluster Comput* 2019;22(Suppl 4):9435–43. <https://doi.org/10.1007/s10586-018-2264-2>
- [284] Wang J, Li T, Yang Y, Chen S, Zhai W. DiagLLM: multimodal reasoning with large language model for explainable bearing fault diagnosis. *Sci China Inf Sci* 2025;68(6):160103.
- [285] Mandelli D, Wang C. A model-based approach to extract health information from textual data. In: 2022 Annual conference of the prognostics and health management society, PHM 2022. Prognostics and Health Management Society; 2022.
- [286] Wang G, Zhang Y, Lu M, Wu Z. Hierarchical graph neural network with adaptive cross-graph fusion for remaining useful life prediction. *Meas Sci Technol* 2023;34(5):055112.
- [287] Li T, Zhao Z, Sun C, Yan R, Chen X. Hierarchical attention graph convolutional network to fuse multi-sensor signals for remaining useful life prediction. *Reliab Eng Syst Saf* 2021;215:107878. <https://doi.org/10.1016/j.res.2021.107878>
- [288] Xiao Y, Liu D, Cui L, Wang H. Heterogeneous graph representation-driven multiplex aggregation graph neural network for remaining useful life prediction of bearings. *Mech Syst Signal Process* 2024;220:111679. <https://doi.org/10.1016/j.ymsp.2024.111679>
- [289] Theiler R, Fink O. Graph neural networks for electric and hydraulic data fusion to enhance short-term forecasting of pumped-storage hydroelectricity. *PHM Society European Conference* 2024;8(1):11. <https://doi.org/10.36001/phme.2024.v8i1.4129>
- [290] Li T, and JD. Multi-level attention graph feature fusion smooth prognostics approach for aircraft engines remaining useful life prediction. *J Control Decis* 2024;0(0):1–17. <https://doi.org/10.1080/23307706.2024.2398536>
- [291] Wang Y, Wu M, Jin R, Li X, Xie L, Chen Z. Local-global correlation fusion-based graph neural network for remaining useful life prediction. *IEEE Trans Neural Netw Learn Syst* 2025;36(1):753–66. <https://doi.org/10.1109/TNNLS.2023.3330487>
- [292] He Y, Su H, Zio E, Peng S, Fan L, Yang Z, et al. A systematic method of remaining useful life estimation based on physics-informed graph neural networks with multisensor data. *Reliab Eng Syst Saf* 2023;237:109333. <https://doi.org/10.1016/j.res.2023.109333>
- [293] Zhang F, Liu J, Li Y, Liu Y, Ge M-F, Jiang X. A health condition assessment and prediction method of Francis turbine units using heterogeneous signal fusion and graph-driven health benchmark model. *Eng Appl Artif Intell* 2023;126:106974. <https://doi.org/10.1016/j.engappai.2023.106974>
- [294] Group U SP. RLfleet: Reinforcement Learning Environment for Fleet Management. <https://github.com/UNSW-MCM/RLfleet>; 2023.
- [295] Leroy P, Morato PG, Pisane J, Kolios A, Ernst D. IMP-MARL: a suite of environments for large-scale infrastructure management planning via MARL. In: *NeurIPS 2023 datasets and benchmarks track*. 2023.
- [296] Hamida Z, Goulet J-A. Hierarchical reinforcement learning for transportation infrastructure maintenance planning. *Reliab Eng Syst Saf* 2023;235:109214. <https://doi.org/10.1016/j.res.2023.109214>
- [297] Hein D, Depeweg S, Tokic M, Udluft S, Hentschel A, Runkler T, et al. A benchmark environment motivated by industrial control problems. In: 2017 IEEE Symposium series on computational intelligence (SSCI). 2017, p. 1–8. <https://doi.org/10.1109/SSCI.2017.8280935>
- [298] Donnot B, et al. Grid2op: a testbed platform to model sequential decision making in power systems 2020; See also: Marot et al., *NeurIPS 2020 for the L2RPN challenge methodology*; <https://github.com/rte-france/Grid2Op>.

- [299] Gamella JL, Peters J, Bühlmann P. Causal chambers as a real-World physical testbed for AI methodology. *Nature Mach Intell* 2025;7:107–18. <https://doi.org/10.1038/s42256-024-00964-x>
- [300] Bhustali P. IMPRL: Inspection and Maintenance Planning using Reinforcement Learning. <https://github.com/omniscientoctopus/implr>; 2023.
- [301] Teubert C, Jarvis K, Corbetta M, Kulkarni C, Daigle M. ProgPy: python packages for prognostics and health management of engineering systems. *J Open Source Softw* 2023;8(87):5099. <https://doi.org/10.21105/joss.05099>
- [302] Raffin A, Hill A, Gleave A, Kanervisto A, Ernestus M, Dormann N. Stable-Baselines3: reliable reinforcement learning implementations. *J Mach Learn Res* 2021;22(268):1–8.
- [303] Brockman G, Cheung V, Pettersson L, Schneider J, Schulman J, Tang J, et al. OpenAI gym. In: Proceedings of the neurIPS deep reinforcement learning workshop. 2016, Standardized API suite for RL environments.
- [304] Torrisi C, et al. Gymnasium: A standard interface for reinforcement learning environments. Farama Foundation, GitHub repository; 2023. Community-maintained successor to OpenAI Gym.
- [305] Sutton RS, Barto AG, et al. Reinforcement learning: An introduction; vol. 1. MIT press Cambridge; 1998.
- [306] Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, et al. Human-level control through deep reinforcement learning. *Nature* 2015;518(7540):529–33.
- [307] Tang C, Abbatemateo B, Hu J, Chandra R, Martín-Martín R, Stone P. Deep reinforcement learning for robotics: a survey of real-world successes. In: Proceedings of the AAAI conference on artificial intelligence; vol. 39. 2025, p. 28694–8.
- [308] Fawzi A, Balog M, Huang A, Hubert T, Romera-Paredes B, Barekatin M, et al. Discovering faster matrix multiplication algorithms with reinforcement learning. *Nature* 2022;610(7930):47–53.
- [309] Silver D, Hubert T, Schrittwieser J, Antonoglou I, Lai M, Guez A, et al. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science* 2018;362(6419):1140–4. <https://doi.org/10.1126/science.aar6404>
- [310] Tian Y, Kladny K-R, Wang Q, Huang Z, Fink O. Multi-agent actor-critic with time dynamical opponent model. *Neurocomputing* 2023;517:165–72.
- [311] Han L, Zhu Q, Sheng J, Zhang C, Li T, Zhang Y, et al. Lifelike agility and play in quadrupedal robots using reinforcement learning and generative pre-trained models. *Nature Mach Intell* 2024;6(7):787–98.
- [312] Esteso A, Peidro D, Mula J, Díaz-Madroño M. Reinforcement learning applied to production planning and control. *Int J Prod Res* 2023;61(16):5772–89.
- [313] Jansen P, Côté M-A, Khot T, Bransom E, Dalvi Mishra B, Majumder BP, et al. Discoveryworld: a virtual environment for developing and evaluating automated scientific discovery agents. *Adv Neural Inf Process Syst* 2024;37:10088–116.
- [314] Lee J, Mitici M. Deep reinforcement learning for predictive aircraft maintenance using probabilistic remaining-useful-life prognostics. *Reliab Eng Syst Saf* 2023;230:108908.
- [315] Zhang H, Wei X, Liu Z, Ding Y, Guan Q. Condition-based maintenance for multi-state systems with prognostic and deep reinforcement learning. *Reliab Eng Syst Saf* 2025;255:110659.
- [316] Lee JS, Yeo I-H, Bae Y. A stochastic track maintenance scheduling model based on deep reinforcement learning approaches. *Reliab Eng Syst Saf* 2024;241:109709.
- [317] Liu H, Li B, Yao F, Hu G, Xie L. Maintenance optimization of multi-unit balanced systems using deep reinforcement learning. *Reliab Eng Syst Saf* 2024;244:109921.
- [318] Knowles M, Baglee D, Wermter S. Reinforcement learning for scheduling of maintenance. In: International conference on innovative techniques and applications of artificial intelligence. Springer; 2010, p. 409–22.
- [319] Barde S RA, Yacout S, Shin H. Optimal preventive maintenance policy based on reinforcement learning of a fleet of military trucks. *J Intell Manuf* 2019;30(1):147–61.
- [320] Rokhforoz P, Montazeri M, Fink O. Safe multi-agent deep reinforcement learning for joint bidding and maintenance scheduling of generation units. *Reliab Eng Syst Saf* 2023;232:109081.
- [321] Chen J, Zhou X. Reinforcement learning based maintenance scheduling of flexible multi-machine manufacturing systems with varying interactive degradation. *Reliab Eng Syst Saf* 2025;260:111018.
- [322] Banerjee C, Nguyen K, Fookes C, Raissi M. A survey on physics informed reinforcement learning: review and open problems. *Expert Syst Appl* 2025;128166.
- [323] Cheng R, Orosz G, Murray RM, Burdick JW. End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control. In: Proceedings of the AAAI conference on artificial intelligence; vol. 33. 2019, p. 3387–95.
- [324] Emam Y, Glotfelter P, Egerstedt M. Safe reinforcement learning using robust control barrier functions. *IEEE Rob Autom Lett* 2021;6(4):6818–25.
- [325] Raissi M, Perdikaris P, Karniadakis GE. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J Comput Phys* 2019;378:686–707.
- [326] Wang R, Cao Z, Zhou X, Wen Y, Tan R. Phyllis: physics-informed lifelong reinforcement learning for data center cooling control. In: Proceedings of the 14th ACM international conference on future energy systems. 2023, p. 114–26.
- [327] Schena L, Marques PA, Poletti R, Ahizi S, Van den Berghe J, Mendez MA. Reinforcement twinning: from digital twins to model-based reinforcement learning. *J Comput Sci* 2024;82:102421.
- [328] Bellani L, Compare M, Baraldi P, Zio E. Towards developing a novel framework for practical phm: a sequential decision problem solved by reinforcement learning and artificial neural networks. *Int J Prognost Health Manage* 2019;10(4).
- [329] Tian Y, Chao MA, Kulkarni C, Goebel K, Fink O. Real-time model calibration with deep reinforcement learning. *Mech Syst Signal Process* 2022;165:108284.
- [330] Zheng Y, Guo R, Liu G. Combining graph neural networks and deep reinforcement learning for aircraft maintenance stand scheduling. In: Proceedings of the 2024 8th international conference on computer science and artificial intelligence. 2024, p. 333–9.
- [331] Liang H, Moya B, Chinesta F, Chatzi E. Resilience-based post disaster recovery optimization for infrastructure system via deep reinforcement learning. *Reliab Eng Syst Saf* 2025;111478.
- [332] Zhang P, Zhu X, Xie M. A model-based reinforcement learning approach for maintenance optimization of degrading systems in a large state space. *Comput Indust Eng* 2021;161:107622.
- [333] Marchetti S, Di Maio F, Zio E. A deep reinforcement learning method for finding the risk-based optimal prescriptive maintenance policy of degrading safety barriers. In: 2024 8th international conference on system reliability and safety (ICRSR). IEEE; 2024, p. 177–83.
- [334] Liang H, Xie J, Huang B, Li Y, Sun B, Yang C. A novel sim2real reinforcement learning algorithm for process control. *Reliab Eng Syst Saf* 2025;254:110639.
- [335] Kumar A, Zhou A, Tucker G, Levine S. Conservative q-learning for offline reinforcement learning. In: Advances in neural information processing systems 33 (NeurIPS 2020). 2020, p. 1179–91.
- [336] Tran QK, Huynh KT, Grall A, Langeron Y. Offline learning of maintenance policies using reinforcement learning and historical maintenance data. In: Proceedings of the 35th european safety and reliability conference (ESREL 2025). 2025, p. 1546–53.
- [337] Chen L, Lu K, Rajeswaran A, Lee K, Grover A, Laskin M, et al. Decision transformer: reinforcement learning via sequence modeling. In: Advances in neural information processing systems 34 (NeurIPS 2021). 2021, p. 15084–97.
- [338] Arora S, Doshi P. A survey of inverse reinforcement learning: challenges, methods and progress. *Artif Intell* 2021;297:103500.
- [339] Deshpande S, Walambe R, Kotecha K, Selvachandran G, Abraham A. Advances and applications in inverse reinforcement learning: a comprehensive review. *Neural Comput Appl* 2025;37:11071–123.
- [340] Banerjee C, Nguyen K, Fookes C, Raissi M. A survey on physics informed reinforcement learning: Review and open problems. 2023. 2309.01909.
- [341] Ramesh A, Ravindran B. Physics-informed model-based reinforcement learning. In: Learning for dynamics and control conference. PMLR; 2023, p. 26–37.
- [342] Xing E, Luk V, Oh J. Stabilizing reinforcement learning in differentiable multi-physics simulation. In: The thirteenth international conference on learning representations. 2025.
- [343] Ma J, Wu F, Chen Y, Ji X, Ding Y. Effective multimodal reinforcement learning with modality alignment and importance enhancement. 2023. 2302.09318.
- [344] Portelas R, Colas C, Weng L, Hofmann K, Oudeyer P-Y. Automatic curriculum learning for deep RL: a short survey. In: Proceedings of the 29th international joint conference on artificial intelligence (IJCAI). 2020.
- [345] Pateria S, Subagdja B, Tan A-h, Quek C. Hierarchical reinforcement learning: a comprehensive survey. *ACM Comput Surv (CSUR)* 2021;54(5):1–35.
- [346] Maggio L GD. Toward autonomous LLM-based AI agents for predictive maintenance. *Appl Sci* 2025;15(21):11515.
- [347] Vyas J, Mercangöz M. Autonomous industrial control using an agentic framework with large language models. *IFAC-PapersOnLine* 2025;59(6):349–54. 14th IFAC Symposium on Dynamics and Control of Process Systems, including Biosystems (DYCOPS 2025). <https://doi.org/10.1016/j.ifacol.2025.07.170>
- [348] Farahani A, Schulte C, et al. Hybrid agentic AI and multi-agent systems in smart manufacturing. In: IEEE international conference on emerging technologies and factory automation (ETFA). 2025, to appear.
- [349] Xie T, Zhao S, et al. Text2reward: reward shaping with large language models. In: ICLR. 2024.
- [350] Liang J, Huang W, Xia F, Xu P, Hausman K, Zeng A. Code as policies: language model programs for embodied control. In: 2023 IEEE International conference on robotics and automation (ICRA). IEEE; 2023, p. 9493–500.
- [351] Lukens S, Brotherton T, Saxena A. Large language model agents as PHM co-pilots. In: PHM Society annual conference. 2024.
- [352] Kirubanandan K, Sankararaman S. Causal-aware LLM agents for PHM co-pilots. In: PHM society annual conference. 2025, to appear.
- [353] Getz N, Tong X. Large language model accelerated maintenance insights. In: PHM society conference. 2025.
- [354] Xia X, et al. A multi-agent convolution deep RL network for aeroengine fleet maintenance optimization. *J Manuf Syst* 2023;68:410–25.
- [355] Sebastián E, Duong T, Atanasov N, Montijano E, Sagüés C. Physics-informed multi-agent reinforcement learning for distributed multi-robot problems. *IEEE Trans Rob* 2025;41:4499–517.
- [356] Rokhforoz P, Montazeri M, Fink O. Multi-agent reinforcement learning with graph convolutional neural networks for optimal bidding strategies of generation units in electricity markets. *Expert Syst Appl* 2023;221:120010. <https://doi.org/10.1016/j.eswa.2023.120010>
- [357] Zintgraf L, Shiarlis K, Igl M, Schulze S, Gal Y, Hofmann K, et al. VariBAD: a very good method for bayes-adaptive deep RL via meta-learning. In: International conference on learning representations. 2020.
- [358] Cheng J, et al. Knowledge transfer for adaptive maintenance policy optimization using meta-reinforcement learning. *Reliab Eng Syst Saf* 2024;247:110127.
- [359] Sebastián E, et al. Physics-informed multi-agent reinforcement learning for distributed robotics. *IEEE Trans Rob* 2025;41:4499–517.
- [360] Michau G, Palmé T, Fink O. Fleet PHM for critical systems: bi-level deep learning approach for fault detection. *PHM Society European Conference* 2018;4(1). <https://doi.org/10.36001/phme.2018.v4i1.403>
- [361] Hu Y, Miao X, Si Y, Pan E, Zio E. Prognostics and health management: a review from the perspectives of design, development and decision. *Reliab Eng Syst Saf* 2022;217:108063. <https://doi.org/10.1016/j.res.2021.108063>

- [362] Cheng J, Cheng M, Liu Y, Wu J, Li W, Frangopol DM. Knowledge transfer for adaptive maintenance policy optimization in engineering fleets based on meta-reinforcement learning. *Reliab Eng Syst Saf* 2024;247:110127. <https://doi.org/10.1016/j.res.2024.110127>
- [363] Feng Y, Chen J, Xie J, Zhang T, Lv H, Pan T. Meta-learning as a promising approach for few-shot cross-domain fault diagnosis: algorithms, applications, and prospects. *Knowl Based Syst* 2022;235:107646. <https://doi.org/10.1016/j.knsys.2021.107646>
- [364] Finn C, Abbeel P, Levine S. Model-agnostic meta-learning for fast adaptation of deep networks. In: *Proceedings of the 34th international conference on machine learning*. PMLR; 2017, p. 1126–35.
- [365] Deleu T, Würfl T, Samiei M, Cohen JP, Bengio Y. Torchmeta: A Meta-Learning library for PyTorch. 2019. <https://github.com/tristandeleu/pytorch-meta>.
- [366] Arnold S MR, Mahajan P, Datta D, Bunner I, Zarkias KS. Learn2learn: a library for meta-learning research 2020;2008.12284.
- [367] Hollmann N, Müller S, Purucker L, Krishnakumar A, Körfer M, Hoo SB, et al. Accurate predictions on small data with a tabular foundation model. *Nature* 2025;637(8045):319–26. <https://doi.org/10.1038/s41586-024-08328-6>
- [368] Grefenstette E, Amos B, Yarats D, Htut PM, Molchanov A, Meier F, et al. Generalized inner loop meta-learning. 2019, arXiv preprint arXiv:191001727 .
- [369] Theiler R, Telyatnikov L, Von Krannichfeldt L, Bizzi A, Fink O. Towards a unified and data-efficient prognostics and health management framework with tabular foundation models; 2025. Manuscript in preparation.
- [370] Gu Z, Zhu B, Zhu G, Chen Y, Tang M, Wang J. Anomalygpt: detecting industrial anomalies using large vision-language models. In: *Proceedings of the AAAI conference on artificial intelligence*; vol. 38. 2024, p. 1932–40.
- [371] Goswami S, Kontolati K, Shields MD, Karniadakis GE. Deep transfer operator learning for partial differential equations under conditional shift. *Nature Machine Intelligence* 2022;4(12):1155–64. <https://doi.org/10.1038/s42256-022-00569-2>
- [372] Song Y, Wang T, Cai P, Mondal SK, Sahoo JP. A comprehensive survey of few-shot learning: evolution, applications, challenges, and opportunities. *ACM Comput Surv* 2023;55(13s):271:1–271:40. <https://doi.org/10.1145/3582688>
- [373] Wang Y, Yao Q, Kwok JT, Ni LM. Generalizing from a few examples: a survey on few-shot learning. *ACM Comput Surv* 2020;53(3):63:1–63:34. <https://doi.org/10.1145/3386252>
- [374] Zhang S, Wei H-L, Ding J. An effective zero-shot learning approach for intelligent fault detection using 1D CNN. *Appl Intell* 2023;53(12):16041–58. <https://doi.org/10.1007/s10489-022-04342-1>
- [375] Xu J, Liang S, Ding X, Yan R. A zero-shot fault semantics learning model for compound fault diagnosis. *Expert Syst Appl* 2023;221:119642. <https://doi.org/10.1016/j.eswa.2023.119642>
- [376] Chen Z, Wu J, Deng C, Wang X, Wang Y. Deep attention relation network: a zero-shot learning method for bearing fault diagnosis under unknown domains. *IEEE Trans Reliab* 2023;72(1):79–89. <https://doi.org/10.1109/TR.2022.3177930>
- [377] Bromley J, Guyon I, LeCun Y, Säckinger E, Shah R. Signature verification using a “siamese” time delay neural network. *Adv Neural Inf Process Syst* 1993;6:737–44.
- [378] Snell J, Swersky K, Zemel R. Prototypical networks for few-shot learning. *Adv Neural Inf Process Syst* 2017;30:4080–90.
- [379] Sung F, Yang Y, Zhang L, Xiang T, Torr P HS, Hospedales TM. Learning to compare: relation network for few-shot learning. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2018, p. 1199–208.
- [380] Hospedales T, Antoniou A, Micaelli P, Storkey A. Meta-learning in neural networks: a survey. *IEEE Trans Pattern Anal Mach Intell* 2022;44(9):5149–69. <https://doi.org/10.1109/TPAMI.2021.3079209>
- [381] Li C, Li S, Feng Y, Gryllias K, Gu F, Pecht M. Small data challenges for intelligent prognostics and health management: a review. *Artif Intell Rev* 2024;57(8):214. <https://doi.org/10.1007/s10462-024-10820-4>
- [382] Finn C, Abbeel P, Levine S. Model-agnostic meta-learning for fast adaptation of deep networks. In: *International conference on machine learning*. PMLR; 2017, p. 1126–35.
- [383] Rajeswaran A, Finn C, Kakade SM, Levine S. Meta-learning with implicit gradients. In: *Advances in neural information processing systems*; vol. 32. Curran Associates, Inc.; 2019.
- [384] Han T, Li Y-F, Qian M. A hybrid generalization network for intelligent fault diagnosis of rotating machinery under unseen working conditions. *IEEE Trans Instrum Meas* 2021;70:1–11. <https://doi.org/10.1109/TIM.2021.3088489>
- [385] Dong Q, Li L, Dai D, Zheng C, Ma J, Li R, et al. A survey on in-context learning. In: Al-Onaizan Y, Bansal M, Chen Y-N, editors. *Proceedings of the 2024 conference on empirical methods in natural language processing*. Miami, Florida, USA: Association for Computational Linguistics; 2024, p. 1107–28. <https://doi.org/10.18653/v1/2024.emnlp-main.64>
- [386] Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, et al. Language models are few-shot learners. In: *Advances in neural information processing systems*; vol. 33. Curran Associates, Inc.; 2020, p. 1877–901.
- [387] Huang Q, Ren H, Chen P, Kržmanc G, Zeng D, Liang PS, et al. Prodigy: enabling in-context learning over graphs. *Adv Neural Inf Process Syst* 2023;36:16302–17.
- [388] Sun Q, Cui Y, Zhang X, Zhang F, Yu Q, Wang Y, et al. Generative multimodal models are in-context learners. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2024, p. 14398–409.
- [389] Taga EO, Ildiz ME, Oymak S. TimePFN: Effective Multivariate Time Series Forecasting with Synthetic Data. 2025. 2502.16294. <https://doi.org/10.48550/arXiv.2502.16294>
- [390] Sun L, Wu J, Wang J, Wen S, Li G, Liu Y. Fault diagnosis of slewing bearing using audible sound signal based on time generative adversarial network-tabPFN method. *J Vib Acoust* 2025;147(4):041002.
- [391] Yang Y, Wang H, Liu Z, Yang Z. Few-shot learning for rolling bearing fault diagnosis via siamese two-dimensional convolutional neural network. In: *2020 11th international conference on prognostics and system health management (PHM-2020 Jinan)*. 2020, p. 373–8. <https://doi.org/10.1109/PHM-Jinan48558.2020.00073>
- [392] Zhao L, Yang Z, Yu Y. Industrial fault diagnosis based on few shot learning. In: *2021 Global reliability and prognostics and health management (PHM-Nanjing)*. 2021, p. 1–6. <https://doi.org/10.1109/PHM-Nanjing52125.2021.9612836>
- [393] Hu Y, Xiong Q, Zhu Q, Yang Z, Zhang Z, Wu D, et al. Few-shot transfer learning with attention for intelligent fault diagnosis of bearing. *J Mech Sci Technol* 2022;36(12):6181–92. <https://doi.org/10.1007/s12206-022-1132-4>
- [394] Wang H, Li C, Ding P, Li S, Li T, Liu C, et al. A novel transformer-based few-shot learning method for intelligent fault diagnosis with noisy labels under varying working conditions. *Reliab Eng Syst Saf* 2024;251:110400. <https://doi.org/10.1016/j.res.2024.110400>
- [395] Costa N, Sánchez L. Few-shot generative compression approach for system health monitoring. *Reliab Eng Syst Saf* 2025;256:110749. <https://doi.org/10.1016/j.res.2024.110749>
- [396] Wang D, Zhang M, Xu Y, Lu W, Yang J, Zhang T. Metric-based meta-learning model for few-shot fault diagnosis under multiple limited data conditions. *Mech Syst Signal Process* 2021;155:107510. <https://doi.org/10.1016/j.ymsp.2020.107510>
- [397] Lei Z, Zhang P, Chen Y, Feng K, Wen G, Liu Z, et al. Prior knowledge-embedded meta-transfer learning for few-shot fault diagnosis under variable operating conditions. *Mech Syst Signal Process* 2023;200:110491. <https://doi.org/10.1016/j.ymsp.2023.110491>
- [398] Zhang T, Chen J, Li F, Zhang K, Lv H, He S, et al. Intelligent fault diagnosis of machines with small & imbalanced data: a state-of-the-art review and possible extensions. *ISA Trans* 2022;119:152–71. <https://doi.org/10.1016/j.isatra.2021.02.042>
- [399] Ding P, Jia M, Zhao X. Meta deep learning based rotating machinery health prognostics toward few-shot prognostics. *Appl Soft Comput* 2021;104:107211. <https://doi.org/10.1016/j.asoc.2021.107211>
- [400] Yang T, Tang T, Wang J, Qiu C, Chen M. A novel cross-domain fault diagnosis method based on model agnostic meta-learning. *Measurement* 2022;199:111564. <https://doi.org/10.1016/j.measurement.2022.111564>
- [401] Mo Y, Li L, Huang B, Li X. Few-shot RUL estimation based on model-agnostic meta-learning. *J Intell Manuf* 2023;34(5):2359–72. <https://doi.org/10.1007/s10845-022-01929-w>
- [402] Yang J, Wang X, Luo Z. Few-shot remaining useful life prediction based on meta-learning with deep sparse kernel network. *Inf Sci* 2024;653. <https://doi.org/10.1016/j.ins.2023.119795>
- [403] Wang D, Xian X, Li H, Wang D. Distribution-agnostic probabilistic few-shot learning for multimodal recognition and prediction. *IEEE Trans Autom Sci Eng* 2024;1–17. <https://doi.org/10.1109/TASE.2024.3372711>
- [404] Li Y, Liu C, Hua J, Gao J, Maropoulos P. A novel method for accurately monitoring and predicting tool wear under varying cutting conditions based on meta-learning. *CIRP Ann* 2019;68(1):487–90. <https://doi.org/10.1016/j.cirp.2019.03.010>
- [405] Yang J, Wang X. Meta-learning with deep flow kernel network for few shot cross-domain remaining useful life prediction. *Reliab Eng Syst Saf* 2024;244. <https://doi.org/10.1016/j.res.2024.109928>
- [406] Ding P, Jia M. Mechatronics equipment performance degradation assessment using limited and unlabeled data. *IEEE Trans Ind Inf* 2022;18(4):2374–85. <https://doi.org/10.1109/TII.2021.3091143>
- [407] Ding P, Xia J, Zhao X, Jia M. Graph structure few-shot prognostics for machinery remaining useful life prediction under variable operating conditions. *Adv Eng Inf* 2024;60:102360. <https://doi.org/10.1016/j.aei.2024.102360>
- [408] Li Y, Wang J, Huang Z, Gao RX. Physics-informed meta learning for machining tool wear prediction. *J Manuf Syst* 2022;62:17–27. <https://doi.org/10.1016/j.jmsy.2021.10.013>
- [409] Majumdar R, Jadhav V, Deodhar A, Karande S, Vig L, Runkana V. HxPINN: a hypernetwork-based physics-informed neural network for real-time monitoring of an industrial heat exchanger. *Numer Heat Transf Part B: Fundament* 2024;0(0):1–22. <https://doi.org/10.1080/10407790.2024.2325648>
- [410] Mu G, Liu C-L, Chen J. Integrating physical knowledge for generative-based zero-shot learning models in process fault diagnosis. *Reliab Eng Syst Saf* 2025;258:110852. <https://doi.org/10.1016/j.res.2025.110852>
- [411] Ma C, Wang X, Li Y, Cai Z. Broad zero-shot diagnosis for rotating machinery with untrained compound faults. *Reliab Eng Syst Saf* 2024;242:109704. <https://doi.org/10.1016/j.res.2023.109704>
- [412] Chen X, Zhao C, Ding J. Pyramid-type zero-shot learning model with multi-granularity hierarchical attributes for industrial fault diagnosis. *Reliab Eng Syst Saf* 2023;240:109591. <https://doi.org/10.1016/j.res.2023.109591>
- [413] Cai L, Deng X, Yin H, Lin J, Qin Y. Generalized zero-sample industrial fault diagnosis with domain bias. *Reliab Eng Syst Saf* 2025;253:110571. <https://doi.org/10.1016/j.res.2024.110571>
- [414] Zhu J, Pang G. Toward generalist anomaly detection via in-context residual learning with few-shot sample prompts. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2024, p. 17826–36.
- [415] Wang S, Wang D, Kong D, Wang J, Li W, Zhou S. Few-shot rolling bearing fault diagnosis with metric-based meta learning. *Sensors* 2020;20(22):6437.
- [416] Vettoruzzo A, Bouguetia M-R, Vanschoren J, Rögnvaldsson T, Santosh KC. Advances and Challenges in Meta-Learning: A Technical Review. 2023. 2307.04722. <https://doi.org/10.48550/arXiv.2307.04722>
- [417] She D, Duan Y, Yang Z, Pecht M. A meta-transfer learning prediction method with few-shot data for the remaining useful life of rolling bearing. *Struct Health Monit* 2025;14759217251321080.

- [418] Song Y, Wang T, Cai P, Mondal SK, Sahoo JP. A comprehensive survey of few-shot learning: evolution, applications, challenges, and opportunities. *ACM Comput Surv* 2023;55(13s):1–40. <https://doi.org/10.1145/3582688>
- [419] Rasul K, Bennett A, Vicente P, Gupta U, Ghonia H, Schneider A, et al. Vq-tr: vector quantized attention for time series forecasting. In: *The twelfth international conference on learning representations*. 2024.
- [420] Li X, Tang H, Chen S, Wang Z, Chen R, Abram M. Why does in-context learning fail sometimes? Evaluating in-context learning on open and closed questions. 2024. 2407.02028. <https://doi.org/10.48550/arXiv.2407.02028>
- [421] Beck J, Vuorio R, Liu EZ, Xiong Z, Zintgraf L, Finn C, et al. A tutorial on meta-reinforcement learning. *Foundat Trends Mach Learn* 2025;18(2-3):224–384.
- [422] Cho W, Lee K, Rim D, Park N. Hypernetwork-based meta-learning for low-rank physics-informed neural networks. *Adv Neural Inf Process Syst* 2023;36:11219–31.
- [423] Bihlo A. Improving physics-informed neural networks with meta-learned optimization. *J Mach Learn Res* 2024;25(14):1–26.
- [424] Vettoruzzo A, Bracciaioli L, Vanschoren J, Nowaczyk M. Unsupervised meta-learning via in-context learning. In: *The thirteenth international conference on learning representations*. 2025.
- [425] Liu R, Zhang Q, Han T, Yang B, Zhang W, Yin S, et al. Survey on foundation models for prognostics and health management in industrial cyber-Physical systems. *IEEE Trans Industr Cyber-Phys Syst* 2024;2:264–80. <https://doi.org/10.1109/TICPS.2024.3425326>
- [426] Yao J, Han T. Utilizing large-scale foundation models for prognostics and health management in wind turbines: techniques, challenges, and future directions. *Renew Sustain Energy Rev* 2026;227:116527. <https://doi.org/10.1016/j.rser.2025.116527>
- [427] He D, He M. Few-shot learning for plastic bearing fault diagnosis – an integrated image processing and NLP approach. *Annual Conference of the PHM Society* 2023;15(1). <https://doi.org/10.36001/phmconf.2023.v15i1.3575>
- [428] Jin M, Wang S, Ma L, Chu Z, Zhang JY, Shi X, et al. Time-LLM: time series forecasting by reprogramming large language models. In: *The twelfth international conference on learning representations*. 2024.
- [429] Gruver N, Finzi M, Qiu S, Wilson AG. Large language models are zero-shot time series forecasters. *Adv Neural Inf Process Syst* 2023;36:19622–35.
- [430] Gao S, Koker T, Queen O, Hartvigsen T, Tsiligkaridis T, Zitnik M. Units: a unified multi-task time series model. *Adv Neural Inf Process Syst* 2024;37:140589–631.
- [431] Yoon H, Tolera BA, Gong T, Lee K, Lee S-J. By my eyes: grounding multimodal large language models with sensor data via visual prompting. In: *EMNLP*. 2024, p. 2219–41.
- [432] Liu Q, Liu X, Liu C, Wen Q, Liang Y. Time-ffm: towards LM-empowered federated foundation model for time series forecasting. *Adv Neural Inf Process Syst* 2024;37:94512–38.
- [433] Jiang K, Gao X, Gao H, Han H, Qi Y, Han H. A zero-shot fault diagnosis framework for chillers based on sentence-level text attributes. *IEEE Trans Ind Inf* 2025;21(4):2878–87. <https://doi.org/10.1109/TII.2024.3514090>
- [434] Liu R, Zhang Q, Han T, Yang B, Zhang W, Yin S, et al. Survey on foundation models for prognostics and health management in industrial cyber-physical systems. *IEEE Trans Industr Cyber-Phys Syst* 2024;2:264–80. <https://doi.org/10.1109/TICPS.2024.3425326>
- [435] Hamann HF, Brunschwiler T, Gjorgiev B, Martins L SA, Puech A, Varbella A, et al. Foundation Models for the Electric Power Grid. 2024. 2407.09434. <https://doi.org/10.48550/arXiv.2407.09434>
- [436] Gulrajani I, Lopez-Paz D. In search of lost domain generalization. 2020, arXiv preprint arXiv:200701434.
- [437] Wang J, Lan C, Liu C, Ouyang Y, Qin T, Lu W, et al. Generalizing to unseen domains: a survey on domain generalization. *IEEE Trans Knowl Data Eng* 2022;1–1. <https://doi.org/10.1109/TKDE.2022.3178128>
- [438] Zhao C, Zio E, Shen W. Domain generalization for cross-domain fault diagnosis: an application-oriented perspective and a benchmark study. *Reliab Eng Syst Saf* 2024;245:109964.
- [439] Lessmeier C, Kimotho JK, Zimmer D, Sextro W. Condition monitoring of bearing damage in electromechanical drive systems by using motor current signals of electric motors: a benchmark data set for data-driven classification. In: *PHM Society european conference*; vol. 3. 2016.
- [440] Li K, Ping X, Wang H, Chen P, Cao Y. Sequential fuzzy diagnosis method for motor roller bearing in variable operating conditions based on vibration analysis. *Sensors* 2013;13(6):8013–41.
- [441] Lu C, Wang Y, Ragulskis M, Cheng Y. Fault diagnosis for rotating machinery: a method based on image processing. *PLoS ONE* 2016;11(10):e0164111.
- [442] Purohit H, Tanabe R, Ichige K, Endo T, Nikaido Y, Suefusa K, et al. Mimi dataset: sound dataset for malfunctioning industrial machine investigation and inspection. In: *Acoustic scenes and events 2019 workshop (DCASE2019)*. 2019, p. 209.
- [443] Li D, Yang Y, Song Y-Z, Hospedales TM. Deeper, broader and artier domain generalization. In: *ICCV*. 2017.
- [444] Venkateswara H, Eusebio J, Chakraborty S, Panchanathan S. Deep hashing network for unsupervised domain adaptation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2017, p. 5018–27.
- [445] Nejjar I, Geissmann F, Zhao M, Taal C, Fink O. Domain adaptation via alignment of operation profile for remaining useful lifetime prediction. *Reliab Eng Syst Saf* 2024;242:109718. <https://doi.org/10.1016/j.rers.2023.109718>
- [446] Xiao Q, Yang M, Yan J, Shi W. Feature decoupling integrated domain generalization network for bearing fault diagnosis under unknown operating conditions. *Sci Rep* 2024;14(1):30848.
- [447] Wang D, Liao B, Zheng J, Huang G, Hua Z, Gu C, et al. Development of regulations, codes and standards on composite tanks for on-board gaseous hydrogen storage. *Int J Hydrogen Energy* 2019;44:22643–53. <https://doi.org/10.1016/j.IJHYDENE.2019.04.133>
- [448] Nejjar I, Frusque G, Forest F, Fink O. Uncertainty-guided alignment for unsupervised domain adaptation in regression. 2024, arXiv preprint arXiv:240113721.
- [449] Bhatt G. Weakly-supervised spatially grounded concept learner for few-shot learning. In: *BMVC* 2023. 2023.
- [450] Ding P, Jia M. Intelligent health evaluation of rolling bearings based on subspace meta-learning. In: *2020 IEEE 18th international conference on industrial informatics (INDIN)*; vol. 1. IEEE; 2020, p. 750–4.
- [451] Dong H, Nejjar I, Sun H, Chatzi E, Fink O. SimMMDG: a simple and effective framework for multi-modal domain generalization. In: *Advances in neural information processing systems (NeurIPS)*. 2023.
- [452] Zhou K, Yang Y, Hospedales T, Xiang T. Deep domain-adversarial image generation for domain generalisation. In: *AAAI*. 2020.
- [453] Nejjar I, Dong H, Fink O. Recall and refine: a simple but effective source-free open-set domain adaptation framework. *Trans Mach Learn Res* 2025.
- [454] Ganin Y, Ustinova E, Ajakan H, Germain P, Larochelle H, Laviolette F, et al. Domain-adversarial training of neural networks. *J Mach Learn Res* 2016;17(59):1–35.
- [455] Li H, Pan SJ, Wang S, Kot AC. Domain generalization with adversarial feature learning. In: *CVPR*. 2018.
- [456] Long M, Cao Y, Wang J, Jordan M. Learning transferable features with deep adaptation networks. In: *International conference on machine learning*. PMLR; 2015, p. 97–105.
- [457] Pan X, Luo P, Shi J, Tang X. Two at once: enhancing learning and generalization capacities via ibn-net. In: *ECCV*. 2018.
- [458] Li D, Yang Y, Song Y-Z, Hospedales TM. Learning to generalize: meta-learning for domain generalization. In: *AAAI*. 2018.
- [459] Huang Z, Wang H, Xing EP, Huang D. Self-challenging improves cross-domain generalization. In: *ECCV*. 2020.
- [460] Dong H, Chatzi E, Fink O. Towards multimodal open-set domain generalization and adaptation through self-supervision. In: *European conference on computer vision*. 2024.
- [461] Xu Q, Zhang R, Zhang Y, Wang Y, Tian Q. A fourier-based framework for domain generalization. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2021, p. 14383–92.
- [462] Lin S, Zhang Z, Huang Z, Lu Y, Lan C, Chu P, et al. Deep frequency filtering for domain generalization. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2023, p. 11797–807.
- [463] Li J, Li Y, Tan J, Liu C. It takes two: dual branch augmentation module for domain generalization. *Neural Netw* 2024;172:106094.
- [464] Chatopadhyay P, Sarangmath K, Vijaykumar V, Hoffman J. Pasta: proportional amplitude spectrum training augmentation for syn-to-real domain generalization. In: *Proceedings of the IEEE/CVF international conference on computer vision (ICCV)*. 2023, p. 19288–300.
- [465] Xu X, Yang J, Shi W, Ding S, Luo L, Liu J. Physaug: a physical-guided and frequency-based data augmentation for single-domain generalized object detection. In: *Proceedings of the AAAI conference on artificial intelligence*; vol. 39. 2025, p. 21815–23.
- [466] Huang Y, Tang B, Yang Q, Ming Z. Physics-informed causal learning network for fault diagnosis of rotating machinery under unseen operating conditions. *Neurocomputing* 2025;639:130187. <https://doi.org/10.1016/j.neucom.2025.130187>
- [467] Xu Z-QJ, Zhang Y, Xiao Y. Training behavior of deep neural network in frequency domain. In: *Neural information processing: 26th international conference, ICONIP 2019, sydney, NSW, australia, december 12–15, 2019, proceedings, Part I* 26. Springer; 2019, p. 264–74.
- [468] Chen Y, Zhang D, Yan R, Xie M. Applications of domain generalization to machine fault diagnosis: a survey. *IEEE/CAA J Autom Sin* 2025;12(1):1963–84.
- [469] Guo Y, Ju G, Zhang J. A domain generalization network for imbalanced machinery fault diagnosis. *Sci Rep* 2024;14(1):25447.
- [470] Shi Y, Deng A, Deng M, Xu M, Liu Y, Ding X, et al. Domain augmentation generalization network for real-time fault diagnosis under unseen working conditions. *Reliab Eng Syst Saf* 2023;235:109188.
- [471] Chu F. Gear pitting fault diagnosis using domain generalizations and specialization techniques. In: *Annual conference of the phm society*. 2023.
- [472] Hong S, Lee J. Multi-label fault diagnosis of rotary machine via domain adversarial neural network-based domain generalization targeting high range rotating speed. In: *PHM society Asia-Pacific conference*. 2023.
- [473] Kim YC, Kim T, Ko JU, Lee J, Kim K. Domain adaptation based fault diagnosis under variable operating conditions of a rock drill. *Int J Prognost Health Manage* 2023;14(2).
- [474] Ren L, Mo T, Cheng X. Meta-learning based domain generalization framework for fault diagnosis with gradient aligning and semantic matching. *IEEE Trans Ind Inf* 2023;20(1):754–64.
- [475] de Oliveira d C PR, Akçay A, Zhang Y, Kaymak U. Remaining useful lifetime prediction via deep domain adaptation. *Reliab Eng Syst Saf* 2020;195:106682.
- [476] Xu J, Ma B, Chen W, Shan C. Asdinnorm: a single-source domain generalization method for the remaining useful life prediction of bearings. *Lubricants* 2024;12(5):175.
- [477] Ding Y, Jia M, Cao Y, Ding P, Zhao X, Lee C-G. Domain generalization via adversarial out-domain augmentation for remaining useful life prediction of bearings under unseen conditions. *Knowl Based Syst* 2023;261:110199.
- [478] Xiao X, Zhang J, Xu D. Contrastive domain-invariant generalization for remaining useful life prediction under diverse conditions and fault modes. *Reliab Eng Syst Saf* 2025;253:110534.

- [479] cheng Wen B, qing Xiao M, qi Wang X, Zhao X, feng Li J, Chen X. Data-driven remaining useful life prediction based on domain adaptation. *PeerJ Comput Sci* 2021;7:e690.
- [480] Wang D, Liu Q, Wu D, Wang L. Meta domain generalization for smart manufacturing: tool wear prediction with small data. *J Manuf Syst* 2022;62:441–9.
- [481] Yan R, Li W, Lu S, Xia M, Chen Z, Zhou Z, et al. Transfer learning for prognostics and health management: advances, challenges, and opportunities. *J Dyn Monit Diagnost* 2024;60–82.
- [482] Dong H, Liu M, Zhou K, Chatzi E, Kannala J, Stachniss C, et al. Advances in multimodal adaptation and generalization: from traditional approaches to foundation models. 2025, *arXiv preprint arXiv:250118592*.
- [483] Dong H, Chatzi E, Fink O. Towards robust multimodal open-set test-time adaptation via adaptive entropy-aware optimization. In: *The thirteenth international conference on learning representations*. 2025.
- [484] Li S, Gong H, Dong H, Yang T, Tu Z, Zhao Y. DPU: dynamic prototype updating for multimodal out-of-distribution detection. In: *Proceedings of the computer vision and pattern recognition conference*. 2025, p. 10193–202.
- [485] Dong H, Zhao Y, Chatzi E, Fink O. Multiood: scaling out-of-distribution detection for multiple modalities. *Adv Neural Inf Process Syst* 2024;37:129250–78.
- [486] Nejjar I, Meunier-Pion J, Frusque G, Fink O. Dg-mix: domain generalization for anomalous sound detection based on self-supervised learning. In: *DCASE*. 2022.