© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723149

Predicting Materials Failure Before It Happens: A Critical Review of AI for Multimodal Degradation Forecasting Under Coupled Environmental Stressors

OPEYEMI OLUWASEGUN LAWAL[1], SAMSON MICHEAL IDOGHOR[2], OGHENEFEGOR FAVOUR UGBINE[3], OPEYEMI OLUWASEGUN LAWAL[4]

[1,4]*Department of Biomedical Engineering, Technische Universität Wien, Wien, Austria, ORCiD: 0009-0001-2746-226X*

[2]*Department of Wood Science and Engineering, Oregon State University, Corvallis, Oregon, USA, ORCiD: 0000-0001-6811-0204*

[3]*Department of Mechanical Engineering, Ecole Centrale de Lyon, Saint-Etienne Campus, France, ORCiD: 0009-0003-6503-5925*

Abstract- Infrastructure, energy, transportation and industrial materials are increasingly exposed to coupled thermal, moisture, chemical and mechanical stressors that can interact to produce nonlinear and history-dependent degradation. Conventional periodic inspection and single-stressor testing often provide limited information about the progression of damage and may fail to capture interactions that influence long-term material performance and failure. Artificial intelligence (AI) provides an opportunity to address these limitations by integrating heterogeneous information from sensors, imaging, experimental measurements, environmental monitoring and physics-based simulations. However, identifying existing damage is fundamentally different from forecasting its future evolution. This critical review examines the transition from AI-based damage detection toward temporal degradation forecasting, remaining useful life estimation and failure probability prediction under coupled environmental and mechanical conditions. Machine learning, deep learning, multimodal learning, physics-informed AI, digital twins, uncertainty quantification and explainability are critically evaluated in relation to their potential and limitations for materials degradation forecasting. Major translational challenges include limited longitudinal datasets, inconsistent experimental protocols, multimodal data alignment, limited cross-material and cross-environment generalization, uncertainty representation, interpretability and the laboratory-to-field gap. In response, an integrated framework is proposed that links environmental exposure and material state evolution with multimodal sensing, AI-based forecasting, physical constraints and uncertainty estimation. This framework provides a pathway toward proactive intervention by connecting predicted degradation trajectories and failure probability with risk-informed inspection and maintenance decisions.

Keywords: Artificial intelligence; materials degradation; degradation forecasting; coupled environmental stressors; multimodal learning; physics-informed AI; digital twins; remaining useful life; uncertainty quantification; failure probability.

I. INTRODUCTION

1.1 Materials Failure as a Multiscale and Multistressor Problem

Materials failure is a complex process that affects infrastructure, energy systems, transportation assets, and industrial components throughout their service life. Degradation rarely results from a single isolated cause. Materials are commonly exposed to changing combinations of temperature, moisture, chemical environments, and mechanical loading, with each stressor capable of altering the response to the others. Thermal cycling can promote expansion, contraction, oxidation, or changes in material properties, while moisture and chemical exposure may accelerate corrosion, hydrolysis, swelling, or other forms of deterioration. Mechanical loading can further initiate and propagate damage through fatigue, cracking, or progressive loss of structural integrity.

These processes occur across multiple spatial and temporal scales, from microstructural and chemical changes to macroscopic damage and eventual functional failure. Importantly, degradation is often nonlinear and path-dependent. The current condition of a material may therefore depend not only on the magnitude of its present exposure, but also on the sequence, duration, and interaction of previous environmental and mechanical conditions. Consequently, failure cannot always be adequately understood from isolated measurements or single-stressor experiments. A meaningful assessment requires consideration of how environmental exposure drives changes in material state over time and how these changes collectively influence the progression toward failure.

1.2 Limitations of Conventional Materials Degradation Assessment

Conventional approaches to materials degradation assessment rely heavily on periodic inspection, accelerated aging tests, single-stressor experiments, and endpoint measurements. Although these methods remain important for characterizing material performance, they provide only limited information about how degradation develops under changing and interacting service conditions. Periodic inspections provide discrete observations rather than a continuous record of material evolution, while accelerated aging experiments may not fully reproduce the sequence, variability, and interaction of stressors encountered during actual service. Similarly, experiments that isolate individual stressors can overlook degradation pathways that emerge from their combined effects.

Maintenance decisions are also frequently guided by predefined thresholds or measurable changes in selected properties. Such approaches can be effective when degradation follows predictable patterns, but they may be less reliable when damage develops gradually, nonlinearly, or through interacting mechanisms. Endpoint testing presents a related limitation because it primarily establishes whether a material has reached a specified condition rather than forecasting how rapidly it will approach that condition. As a result, conventional assessment can identify visible or measurable damage only after substantial degradation has already occurred, reducing the opportunity for timely intervention and limiting the ability to anticipate future material failure.

1.3 Emergence of AI for Materials Degradation Forecasting

The development of artificial intelligence has created new opportunities to move materials assessment beyond conventional inspection and condition-based evaluation. Machine learning and deep learning can identify relationships between environmental exposure, material properties, and observed degradation, while computer vision can extract damage-related information from images and other visual measurements, as demonstrated for crack identification in concrete and steel surfaces (Cha et al., 2017) and for corrosion recognition from surface imagery (Atha & Jahanshahi, 2018; Nash et al., 2022). Sensor fusion further enables information from multiple monitoring systems to be considered together, potentially providing a more comprehensive representation of material condition. Physics-informed AI and digital twins extend this capability by incorporating physical knowledge, simulations, and continuously updated representations of material behaviour into predictive frameworks (Karpatne et al., 2017; Wang et al., 2025).

However, the increasing use of AI does not by itself establish the ability to forecast failure. An important distinction must be made between recognising damage that already exists and predicting how that damage will evolve. A model that accurately identifies cracks, corrosion, defects, or other manifestations of degradation is primarily performing a detection or classification task. Forecasting requires a different capability: estimating the future trajectory of material degradation from its current state, exposure history, and evolving operating conditions. This distinction is particularly important under coupled stressors, where changes in one condition may alter the rate or mechanism of degradation over time. Consequently, the value of AI for materials failure assessment should be judged not only by its ability to recognise existing damage, but also by whether it can provide reliable and sufficiently early predictions of future material deterioration.

1.4 Scope and Critical Questions

This review examines the extent to which artificial intelligence can support reliable forecasting of materials degradation under coupled environmental and mechanical stressors. Particular attention is given to the integration of complementary data modalities, including sensor measurements, imaging, experimental observations, environmental exposure records, and physics-based simulation data. The review considers whether combining these sources can provide a more informative representation of evolving material condition than any individual modality alone.

Beyond predictive accuracy, the review evaluates whether incorporating physical knowledge and explicit uncertainty estimation can improve the reliability of AI-based forecasts. This is particularly important when models are required to predict degradation beyond the conditions represented in their training data, where apparently accurate predictions may not remain physically meaningful. The analysis therefore considers the role of physics-informed approaches, uncertainty quantification, and explainability in establishing confidence in predicted degradation trajectories and failure probabilities.

A further focus is the gap between controlled laboratory studies and real-world material behaviour. Differences in environmental exposure, loading histories, material variability, measurement quality, and available failure data can substantially affect model performance outside the conditions under which it was developed. Accordingly, the review asks four central questions: how effectively can AI forecast degradation under coupled stressors, which data modalities provide genuinely complementary information, whether physical constraints and uncertainty estimation improve forecasting reliability, and what factors continue to limit translation from laboratory demonstrations to field applications.

II. COUPLED ENVIRONMENTAL STRESSORS AND THE PHYSICS OF MATERIALS DEGRADATION

2.1 Thermal Stress and Thermally Driven Degradation

Temperature is a major driver of materials degradation because changes in thermal conditions can alter both the physical state and mechanical response of a material. Repeated heating and cooling produce thermal cycling, causing expansion and contraction that can generate internal stresses, particularly when adjacent materials have different coefficients of thermal expansion. Expansion mismatch at interfaces can therefore contribute to progressive damage during repeated temperature changes.

Sustained exposure to elevated temperatures can also promote time-dependent deformation such as creep, while thermal conditions may accelerate oxidation and other temperature-dependent chemical processes. In some materials, sufficiently high or changing temperatures can induce phase transformations that alter microstructure and consequently affect mechanical and functional properties. These mechanisms do not necessarily develop independently. Their progression depends on the magnitude and duration of thermal exposure, the frequency of temperature changes, material composition, and the prior thermal history. Thermal degradation is therefore better understood as an evolving process in which temperature exposure progressively modifies the material state and its susceptibility to subsequent damage.

2.2 Moisture and Chemical Exposure

Moisture and chemical exposure can substantially alter the condition and performance of materials by initiating or accelerating physical and chemical degradation processes. Water ingress through pores, surfaces, interfaces, or existing defects can change the internal environment of a material and promote processes such as hydrolysis, corrosion, and swelling. The effects depend on the material structure and composition, as well as the duration and conditions of exposure.

Chemical exposure can similarly cause deterioration through reactions between aggressive species and the material. Chemical attack may alter the material structure, reduce its mechanical integrity, or change its surface and internal properties. Moisture can further influence these processes by facilitating the transport of reactive species and modifying local chemical conditions. In reinforced concrete, for example, carbonation and chloride ingress progressively reduce the alkalinity that protects embedded steel, initiating corrosion that machine learning models increasingly aim to predict from exposure and mix-design variables (Fuhaid & Niaz, 2022; Jia et al., 2022). In polymers,

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723149

moisture-driven hydrolysis and chain scission similarly couple with thermal and mechanical exposure to govern aging behaviour (Zeng et al., 2025). As a result, moisture and chemical degradation may evolve together rather than as independent mechanisms. Their effects can also depend on previous exposure and the condition of the material, making degradation difficult to describe using exposure level alone. Understanding these processes therefore requires consideration of both the external environment and the changing internal state of the material.

2.3 Mechanical Loading and Fatigue

Mechanical loading contributes to materials degradation through the gradual accumulation of damage under repeated or cyclic stresses. Unlike a single loading event, cyclic loading can produce progressive changes within a material even when individual loading cycles remain below the level required for immediate failure. Over time, repeated stress can lead to the initiation of cracks at susceptible locations such as defects, interfaces, or regions of local stress concentration.

Once initiated, cracks may propagate with continued loading, reducing the effective load-bearing capacity of the material and increasing its susceptibility to further damage. The rate of fatigue accumulation and crack growth depends on factors including the applied stress, loading history, material condition, and environmental exposure (Hamada et al., 2025). Mechanical degradation is therefore not determined solely by the current loading condition. Previous loading cycles can influence the material state and affect its subsequent response. This cumulative and history-dependent behaviour makes fatigue particularly important when forecasting degradation over extended service periods.

2.4 Synergistic and Antagonistic Stressor Interactions

The effects of environmental and mechanical stressors are not necessarily additive when they occur simultaneously. A material exposed to temperature, moisture, chemical agents, and mechanical loading may experience degradation that differs substantially from the response expected from each stressor considered independently. In some cases, one stressor can increase the material's susceptibility to another, producing a synergistic effect and accelerating degradation. For example, environmental exposure may alter the material state in a way that makes it more vulnerable to subsequent mechanical damage.

Interactions can also be antagonistic when one process limits, modifies, or competes with another degradation mechanism. Such competing effects make the overall response difficult to predict from isolated stressor measurements. The resulting degradation may therefore exhibit nonlinear behaviour, where relatively small changes in exposure or loading conditions produce disproportionately large changes in damage progression. These interactions also mean that degradation history is important because the effect of a current stressor may depend on changes produced by previous exposure. Consequently, realistic degradation forecasting requires consideration of interacting stressors and their evolving effects rather than treating each degradation mechanism as an independent contribution.

2.5 Environmental Exposure to Material State Evolution

Materials degradation can be understood as a progression from external environmental exposure to changes within the material, followed by observable damage and eventual functional failure. Temperature, moisture, chemical conditions, and mechanical loading influence the internal material state by altering its microstructure, chemical condition, mechanical properties, or accumulated damage. These changes may initially remain difficult to observe, but can subsequently manifest as measurable features such as cracking, corrosion, deformation, surface deterioration, or loss of mechanical performance.

The relationship between exposure and failure is therefore indirect and evolves over time. The same environmental condition may produce different outcomes depending on the material's initial condition and its previous exposure history. Similarly, observable damage represents only one stage in a broader degradation process and may not fully reflect the underlying material state. This progression highlights the importance of monitoring both environmental conditions and material responses when developing forecasting approaches. For AI-based degradation prediction, the objective is consequently not simply to associate environmental variables with visible damage, but to learn how changing exposure conditions influence material state and how that evolving state contributes to future functional failure. The major degradation pathways discussed above produce measurable changes in material condition that can be captured through different experimental and monitoring approaches. Table 1 summarizes representative material classes, their dominant degradation mechanisms, the stressors that contribute to deterioration, the observable signals that can be used for monitoring, and AI approaches relevant to degradation assessment and forecasting. The table provides a concise link between the physical degradation process and the types of data that can support AI-based prediction.

Table 1. Materials, Degradation Mechanisms, Stressors, Observable Signals and Relevant AI Approaches

Material class	Main degradation mechanisms	Key stressors	Observable signals	Relevant AI approaches
Metals and alloys	Corrosion, fatigue, creep, oxidation	Temperature, moisture, chemicals, cyclic loading	Cracks, corrosion features, strain, vibration, acoustic emission	CNNs, LSTM, random forests, probabilistic models
Polymers	Hydrolysis, oxidation, swelling, embrittlement	Moisture, temperature, chemicals, mechanical loading	Mass change, surface damage, stiffness, thermal response	Regression models, CNNs, LSTM, transformers
Concrete and cementitious materials	Cracking, corrosion, carbonation, moisture-related deterioration	Moisture, temperature, chemicals, mechanical loading	Cracks, strain, acoustic emission, stiffness loss	Computer vision, CNNs, random forests, temporal models
Fiber-reinforced composites	Matrix cracking, delamination, fatigue, fiber damage	Temperature, moisture, cyclic loading	Delamination, strain, acoustic emission, stiffness loss	CNNs, LSTM, multimodal models
Protective coatings	Cracking, delamination, corrosion-related deterioration	Moisture, chemicals, temperature, mechanical damage	Surface defects, thickness changes, adhesion loss	Computer vision, CNNs, anomaly detection

III. FROM DAMAGE DETECTION TO DEGRADATION FORECASTING: THE AI TRANSITION

3.1 Conventional AI for Damage Detection

The early application of artificial intelligence to materials assessment has focused largely on the automated recognition of damage that is already present. Machine learning and deep learning methods have been used to identify features associated with cracks, corrosion, surface defects, and other forms of material deterioration. Computer vision has been particularly useful for image-based assessment: convolutional neural networks trained on concrete and steel surface imagery have achieved high accuracy in crack classification (Cha et al., 2017), and comparable architectures have been applied to corrosion recognition, including approaches that report calibrated confidence alongside their predictions (Atha & Jahanshahi, 2018; Nash et al., 2022).

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723149

These approaches can improve the consistency and efficiency of damage assessment by processing large numbers of observations and identifying patterns within complex datasets. Convolutional neural networks and related deep learning methods can learn visual features directly from images, while conventional machine learning approaches can classify damage using selected measurements or extracted features. However, these systems primarily address the question of whether a particular form of damage can be identified from available observations. Their output therefore describes the present or previously observed condition of the material rather than necessarily indicating how that condition will change with continued exposure or loading.

This distinction is important because reliable materials management requires more than recognising damage after it has developed. Detection provides information about the current material state, whereas forecasting requires an understanding of its temporal evolution and the conditions that may influence future degradation. The transition from automated damage recognition to degradation forecasting therefore represents a fundamental shift in the role of AI within materials failure assessment.

3.2 Why Detection Is Not Prediction

Damage detection and degradation forecasting address fundamentally different questions. Detection asks what damage is present at a given point in time, such as whether a crack, corrosion feature, or surface defect can be identified. Forecasting asks how that damage is likely to evolve under continued exposure and loading, and when the resulting degradation may make functional failure probable. A model can therefore perform well in recognising existing damage without demonstrating that it can predict its future progression.

This distinction is particularly important because degradation is influenced by material history and changing environmental and mechanical conditions. The presence of a crack, for example, provides evidence of an existing damage state, but does not by itself determine its future growth. Forecasting requires information about the factors governing progression and their evolution over time. Consequently, moving from detection to prediction requires models that capture temporal changes in material condition rather than relying only on static observations.

3.3 Temporal Degradation Modelling

Temporal degradation modelling addresses the changing condition of a material by incorporating observations collected over time. Time-series models can represent relationships between sequential measurements and identify patterns associated with progressive changes in material behaviour. Recurrent neural networks, including long short-term memory (LSTM) models, have been used to represent temporal dependencies through gated memory cells that mitigate the vanishing-gradient problem of earlier recurrent architectures (Hochreiter & Schmidhuber, 1997), while temporal convolutional models provide another approach for extracting patterns across sequential observations.

More recently, transformer-based architectures, which rely on self-attention rather than recurrence to relate observations across a sequence (Vaswani et al., 2017), have provided mechanisms for modelling relationships across longer sequences and can also support the integration of information from different data sources. Their application to degradation forecasting, however, depends on the availability of sufficiently informative longitudinal data. The ability of a temporal model to reproduce patterns within observed sequences does not necessarily demonstrate that it can reliably predict degradation beyond the conditions represented in those observations. Temporal modelling therefore provides an important foundation for forecasting, but its reliability remains dependent on data quality, exposure history, and the physical relevance of the information used by the model.

3.4 Forecasting Degradation Trajectories

AI-based degradation forecasting extends materials assessment from identifying existing damage to estimating how material condition will change over time. Rather than producing a single classification or condition label, forecasting models can learn relationships between historical observations, environmental exposure and mechanical loading to estimate future degradation trajectories. These trajectories may describe crack growth, corrosion progression, strength loss, fatigue accumulation or changes in other material properties (Hamada et al., 2025). Such predictions are particularly valuable when they provide information about the expected rate and direction of degradation rather than simply identifying whether damage is present.

A further objective is to translate predicted degradation trajectories into engineering outcomes such as remaining useful life and failure probability. Remaining useful life estimation seeks to determine the expected time until a material or component reaches a defined performance or failure condition, while failure probability provides an estimate of the likelihood that failure will occur within a specified future period. These outputs can support condition-based maintenance and risk-informed decision making (Osinubi et al., 2026). However, their reliability depends on whether the model adequately represents the temporal evolution of degradation and the changing conditions under which that evolution occurs. Forecasting therefore requires models that can capture both the observed state of a material and the processes governing its progression toward functional failure.

3.5 Critical Limitation

Although AI models may achieve high predictive accuracy over short forecasting horizons, such performance does not necessarily indicate reliable long-horizon prediction of materials degradation. Degradation processes can evolve through nonlinear, history-dependent and changing mechanisms, meaning that small prediction errors can accumulate as the forecasting horizon increases. A model trained on relatively short sequences may therefore perform well when predicting the near future while becoming increasingly uncertain when extrapolating beyond the temporal conditions represented in its training data.

Long-horizon forecasting is further complicated by changes in environmental exposure, mechanical loading and material state that may not follow patterns observed during model development. Consequently, short-term predictive accuracy should not be treated as sufficient evidence that an AI system can reliably anticipate long-term degradation or failure. Robust forecasting requires evaluation across extended time periods and changing operating conditions, together with appropriate treatment of uncertainty and model generalization.

IV. MULTIMODAL AI FOR MATERIALS DEGRADATION

4.1 Why a Single Data Modality Is Insufficient

Materials degradation is a multidimensional process that produces changes across physical, chemical, mechanical, thermal and morphological domains. These changes do not necessarily occur at the same rate or remain directly observable through a single measurement technique. For example, environmental exposure may alter temperature and moisture conditions while simultaneously driving chemical reactions, changes in microstructure, mechanical property loss and the development of visible or subsurface defects. A single data modality can therefore capture only part of the evolving degradation state.

This limitation is particularly important when degradation results from coupled stressors. Sensor measurements may provide continuous information about environmental and mechanical conditions, while imaging can reveal changes in surface morphology or internal structure. Mechanical measurements can indicate changes in stiffness, strength or fatigue behaviour, whereas chemical measurements may provide evidence of reactions occurring before substantial physical damage becomes observable. Considering these sources together can provide a more complete representation of both the material's current condition and its evolution over time.

Consequently, multimodal AI provides an opportunity to integrate complementary observations rather than relying on one measurement stream (Baltrušaitis et al., 2019). Combining different modalities can help models distinguish environmental exposure from material response and relate observable changes to underlying degradation processes. However, the value of multimodal data depends not simply on the number of measurements available, but on whether the different data sources are physically relevant, temporally aligned and sufficiently reliable for forecasting.

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723149

4.2 Major Data Modalities

Multimodal materials degradation forecasting can draw on several complementary data sources. Sensor and time-series data provide continuous measurements of temperature, humidity, strain, vibration, acoustic emission and chemical signals, allowing changes in environmental exposure and material response to be tracked over time. Imaging methods, including optical imaging, microscopy, thermal imaging, X-ray or CT and, where appropriate, hyperspectral imaging, provide spatial information about surface and internal changes that may not be captured by sensors alone.

Experimental and mechanical data, such as stress–strain behaviour, fatigue measurements and fracture characteristics, provide direct evidence of changes in material performance. Environmental data describe exposure history, cycling conditions and loading conditions that influence degradation. Simulation and physics-based data, including finite element and other mechanistic models, can further provide information about material behaviour under conditions that may be difficult to measure experimentally. Together, these modalities provide complementary information for representing both material state and its evolution.

4.3 Multimodal Fusion Strategies

Multimodal fusion determines how information from different data sources is combined before or during model training. Following the taxonomy proposed by Baltrušaitis et al. (2019), in early fusion, measurements from different modalities are combined at the input stage, providing the model with a joint representation of the available observations. Feature-level fusion first extracts relevant features from each modality and then combines these representations for prediction. Late fusion keeps the modalities separate for most of the modelling process and combines their individual predictions at a later stage. These approaches differ in complexity and in how effectively they preserve information that is specific to each modality.

More advanced approaches use cross-modal attention to allow the model to learn relationships between different data sources, such as linking environmental conditions with changes observed in images or sensor signals; this mechanism builds on the self-attention formulation originally introduced for sequence modelling (Vaswani et al., 2017). Multimodal transformers extend this approach by modelling relationships across multiple sequences or representations within a common forecasting framework. Their usefulness, however, depends on the quality and compatibility of the input data. Fusion does not automatically improve prediction; poorly aligned or weakly informative modalities can introduce noise rather than useful information.

4.4 The Data Alignment Problem

Integrating multimodal data is challenging because different measurements are collected at different sampling rates and temporal or spatial scales. Environmental sensors may provide continuous measurements, while imaging, mechanical testing and other experimental measurements may be collected intermittently. Missing observations, sensor drift and measurement noise can further complicate the alignment of these datasets and make it difficult to establish reliable relationships between environmental exposure and material response.

For degradation forecasting, these alignment issues are particularly important because the timing of an observation can influence its interpretation. A structural change observed in an image, for example, must be related appropriately to the environmental and loading conditions that preceded it. Poor temporal or spatial alignment can therefore introduce errors into multimodal models and may lead to misleading associations between different measurements. Reliable fusion consequently requires careful handling of sampling differences, missing data, sensor stability and measurement quality.

4.5 Does More Data Actually Produce Better Forecasts?

A larger volume of data does not necessarily lead to better materials degradation forecasts. The usefulness of data depends primarily on its physical relevance, consistency and representativeness. Large datasets collected under limited or unrealistic exposure conditions may provide substantial statistical information while offering little insight into how materials will degrade under actual service environments. Similarly, inconsistent measurements,

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723149

differences between testing procedures and sensor errors can introduce uncertainty that additional data cannot simply eliminate.

Dataset bias presents a further limitation. If the available data overrepresent particular materials, environmental conditions, degradation mechanisms or failure modes, an AI model may learn patterns that do not generalize to conditions outside the training dataset. Consequently, improving degradation forecasting requires not only increasing data volume but also ensuring that datasets capture relevant material states, exposure histories and degradation pathways with sufficient quality and diversity.

V. AI ARCHITECTURES FOR FORECASTING DEGRADATION AND FAILURE

5.1 Classical Machine Learning

Classical machine learning methods remain important for materials degradation forecasting, particularly where experimental datasets are limited in size. Random forests, which aggregate predictions across an ensemble of decision trees (Breiman, 2001), and gradient-boosting methods such as XGBoost (Chen & Guestrin, 2016) can capture nonlinear relationships between environmental conditions, material characteristics and measured degradation responses, while support vector machines are effective for regression and classification when the number of variables is relatively large compared with the number of observations. Gaussian processes (Rasmussen & Williams, 2006) are particularly valuable because they provide both predictions and an estimate of predictive uncertainty, which is important when forecasting degradation under limited data.

Despite these advantages, classical methods generally rely on predefined features and therefore may have difficulty representing complex temporal behaviour directly from sequential observations. Their performance can also depend strongly on feature selection and the quality of the available measurements. Thus, while they can provide robust solutions for smaller datasets, they may be less suitable for long-term degradation processes involving multiple interacting variables and evolving material states.

5.2 Deep Learning

Deep learning methods are increasingly used for materials degradation forecasting because they can learn complex representations directly from high-dimensional data, a capability whose modern form was established by early large-scale image classification networks (Krizhevsky et al., 2012) and later consolidated as a general paradigm across scientific and engineering domains (LeCun et al., 2015). Convolutional neural networks (CNNs) are particularly suited to imaging data, where they can extract spatial features associated with cracks, corrosion and other structural changes (Cha et al., 2017; Nash et al., 2022). For sequential measurements, recurrent neural networks (RNNs), including long short-term memory (LSTM) models (Hochreiter & Schmidhuber, 1997), can learn temporal patterns from sensor and degradation data collected over time.

Autoencoders, which learn compact low-dimensional codes by training a network to reconstruct its own input (Hinton & Salakhutdinov, 2006), provide another approach by learning compact representations of complex material data and can support feature extraction and identification of unusual degradation behaviour. Deep regression models can also map multiple input variables directly to continuous degradation measures, such as changes in material properties or predicted damage levels. However, the effectiveness of these approaches depends strongly on the availability of sufficiently large, representative and consistent datasets.

5.3 Transformer-Based Temporal Forecasting

Transformer-based models, built on the self-attention mechanism introduced by Vaswani et al. (2017), offer an effective approach to degradation forecasting because attention can capture relationships between observations separated by long time intervals. This is particularly relevant for materials whose degradation depends on accumulated exposure, loading history and changes in material state. Transformers can also integrate different data types within a common framework, making them suitable for multimodal forecasting involving sensor measurements,

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723149

environmental conditions, imaging features and other material information.

Despite these advantages, transformer-based models have important limitations. They generally require substantial amounts of training data and can involve considerable computational demands, which may be difficult to justify for small materials datasets. Their predictions can also be difficult to interpret because the relationships learned by the model may not correspond directly to known degradation mechanisms. Therefore, their value for materials forecasting should be assessed not only by predictive accuracy but also by data requirements, computational feasibility and the physical interpretability of their predictions.

5.4 Probabilistic and Survival-Based Models

Probabilistic and survival-based models are useful when degradation forecasting involves uncertainty in both the timing and occurrence of failure. Rather than predicting a single time to failure, these approaches can estimate the probability that a material or component will fail within a specified period. Survival analysis, originating from methods such as the proportional-hazards framework (Cox, 1972), provides a basis for modelling time to failure, while hazard models describe how the risk of failure changes as degradation progresses and operating conditions vary.

These approaches are also relevant to remaining useful life estimation, where the objective is to determine how long a material or component is expected to continue operating before reaching a defined failure condition. Their probabilistic formulation is particularly valuable for materials degradation because failure rarely occurs at an identical time under variable exposure and loading histories. By representing failure risk and time to failure explicitly, probabilistic models can provide information that is more useful for maintenance and risk-based decision making than a single deterministic forecast.

VI. PHYSICS-INFORMED AI: CONSTRAINING PREDICTION WITH MATERIALS SCIENCE

6.1 Why Purely Data-Driven Models Are Insufficient

Purely data-driven AI models can provide strong predictions within the conditions represented in their training data, but their reliability can decline when they are required to extrapolate beyond those conditions. Materials degradation often involves physical processes that change with exposure, loading history and material state. A model that learns statistical relationships without representing these underlying processes may therefore produce predictions that are mathematically plausible but physically unrealistic, particularly over long forecasting horizons.

The problem is further complicated by sparse failure data and distribution shifts between training and deployment conditions. Actual failure events are relatively uncommon, while laboratory datasets may differ substantially from field conditions in terms of materials, environments and loading histories. When the operating distribution changes, relationships learned from the original dataset may no longer remain valid. These limitations indicate the need to incorporate physical knowledge into AI forecasting so that predictions remain consistent with known degradation behaviour and are more robust when data are limited or conditions change (Karpatne et al., 2017).

6.2 Physics-Informed Neural Networks and Hybrid Models

Physics-informed neural networks (PINNs) incorporate established physical equations and constraints into the training of neural networks by penalising deviations from governing partial differential equations during learning (Raissi et al., 2019). Instead of relying entirely on observed data, the model is guided by known relationships that describe material behaviour, such as conservation laws, transport processes or degradation kinetics. This can help constrain predictions and reduce physically unrealistic behaviour, particularly when experimental data are limited.

Hybrid models take a broader approach by combining data-driven learning with established physics-based models (Karniadakis et al., 2021). The AI component can learn relationships that are difficult to describe explicitly, while the physical component provides constraints based on known degradation mechanisms. This combination can improve the physical consistency of forecasts and provide a more appropriate framework for predicting degradation under conditions that are not fully represented in the available data.

6.3 Degradation Kinetics and Constitutive Relationships

Degradation kinetics and constitutive relationships provide a physical basis for describing how material properties change under environmental and mechanical conditions. Reaction kinetics can represent the rate of chemical degradation, while diffusion relationships describe the movement of moisture or reactive species through a material. Heat transfer relationships are similarly important for representing temperature changes and their influence on degradation processes.

For mechanically driven degradation, crack growth relationships and fatigue laws can describe the evolution of defects under repeated loading. Incorporating these relationships into AI models provides information about the mechanisms governing degradation rather than relying solely on statistical patterns in observed data (Hamada et al., 2025). This can help constrain forecasts and improve their physical consistency, particularly when predicting material behaviour under conditions that differ from those directly represented in the training data.

6.4 Hybrid Physics–AI Architectures

Hybrid physics–AI architectures combine physical knowledge with machine learning to improve the reliability of degradation predictions. Physics-guided machine learning incorporates known physical relationships into model development (Karpatne et al., 2017), while physics-informed neural networks constrain learning using governing equations or physical conditions (Raissi et al., 2019). Mechanistic machine learning hybrids similarly combine established degradation models with data-driven components, allowing each approach to address limitations of the other.

Surrogate modelling provides another useful strategy by using machine learning to approximate computationally expensive physics-based simulations. This can substantially reduce the computational cost of repeated predictions while retaining information derived from mechanistic models. These approaches are particularly relevant when degradation data are limited or when predictions must remain consistent with known material behaviour. However, their effectiveness depends on the validity of the underlying physical model and the quality of its integration with the learning framework (Karniadakis et al., 2021).

6.5 Remaining Challenge

Although physics-informed and hybrid approaches can improve the physical consistency of AI predictions, the physical models themselves are not always complete or universally applicable. Degradation mechanisms can involve interacting processes that are difficult to represent through a single set of equations or constitutive relationships. Physical models may also contain uncertain parameters and assumptions that are valid only within particular temperature, loading, environmental or material regimes.

This creates an important limitation when physics-based models are incorporated into AI systems. If the underlying physical description is inaccurate or becomes invalid under new conditions, the resulting AI forecast may inherit these limitations. Therefore, physics-informed modelling should not be viewed as a guarantee of reliable prediction. Its effectiveness depends on the quality, uncertainty and range of validity of the physical knowledge incorporated into the forecasting framework.

VII. DIGITAL TWINS AND CONTINUOUS MATERIALS-STATE FORECASTING

7.1 From AI Models to Digital Representations

Digital twins extend conventional AI models by creating a dynamic digital representation of a physical material, component or structure throughout its service life. The concept, originally formulated as a way of

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723149

linking a physical asset to its virtual equivalent to mitigate unpredictable emergent behaviour across the product lifecycle (Grieves & Vickers, 2017), was among the earliest applied to structural life prediction in aerospace, where digital twins were proposed to combine high-fidelity vehicle models with sensor data to forecast fatigue life at the level of an individual airframe (Tuegel et al., 2011). Instead of analysing individual measurements independently, the digital representation can combine information on the material's current state, environmental exposure, loading history and structural response. This allows changes in material condition to be tracked continuously and provides a basis for relating observed degradation to its previous exposure and evolving state.

For materials degradation, this distinction is important because the condition of a material at a given time reflects its accumulated history. A digital twin can therefore maintain an evolving representation of properties such as stiffness, strength, damage state or crack condition as new observations become available. The objective is not simply to reproduce the physical system, but to maintain a useful representation that can support estimation of the present state and prediction of its future evolution.

7.2 AI-Enabled Digital Twins

AI-enabled digital twins integrate sensor observations, physics-based simulations and AI forecasting within a continuously updated framework (Wang et al., 2025). Sensor data provide information about the actual operating and environmental conditions, while physics-based models provide knowledge of material behaviour and degradation mechanisms. AI models can then identify patterns in the collected data and forecast changes that may not be readily captured by mechanistic models alone.

Continuous updating is a central feature of this approach. As new measurements become available, the estimated state of the digital twin can be revised and future degradation forecasts can be updated accordingly. This allows predictions to respond to changing environmental conditions, loading histories and material states rather than remaining fixed after initial model development. Such integration provides a pathway from conventional monitoring toward continuous prediction of material condition.

7.3 Digital Twins for Predictive Maintenance

Digital twins can support predictive maintenance by providing estimates of remaining useful life and changes in failure risk as material degradation progresses (Osinubi et al., 2026; Wang et al., 2025). Instead of relying exclusively on fixed inspection intervals or predefined maintenance thresholds, maintenance decisions can be informed by the predicted condition and future behaviour of the component. This can help identify situations in which intervention may be required before degradation reaches a critical level.

Another advantage is the ability to evaluate alternative intervention and operating scenarios. A digital twin can be used to examine how changes in loading, environmental exposure or maintenance timing may influence future material condition. This supports risk-based intervention and allows maintenance strategies to be compared using predicted consequences rather than relying solely on the current observed state.

7.4 Critical Barriers

Despite their potential, digital twins face several practical limitations. They require reliable data over extended periods, and sensor failures, measurement noise, drift and missing observations can reduce the accuracy of the estimated material state. Continuous model updating also requires methods that can incorporate new information without producing unstable or inconsistent predictions. In addition, complex physics-based simulations and AI models may require substantial computational resources, particularly when digital twins are expected to operate continuously (Wang et al., 2025).

Data integrity and cybersecurity are also important considerations because digital twins depend on continuous information exchange between physical systems, sensors, computational models and decision platforms. Incorrect, corrupted or unreliable data can therefore affect both state estimation and subsequent forecasts. For digital twins to provide dependable materials degradation forecasting, improvements in sensing, data management, model updating, computational efficiency and data security must develop alongside advances in AI modelling.

VIII. UNCERTAINTY QUANTIFICATION AND EXPLAINABILITY

8.1 Why Accuracy Alone Is Insufficient

Predictive accuracy is an important measure of AI performance, but it is not sufficient for establishing the reliability of materials failure forecasts. A model may produce accurate predictions on a test dataset while performing poorly when applied to different materials, exposure conditions or stages of degradation. This is particularly important for failure forecasting because the consequences of an incorrect prediction can be greater as the forecast moves further into the future. A prediction of remaining useful life, for example, is more useful when accompanied by an indication of how certain that estimate is.

A trustworthy forecasting system should therefore communicate both its expected outcome and the confidence associated with that outcome. This allows users to distinguish between predictions supported by substantial evidence and those produced under conditions where the model has limited experience. Accuracy should consequently be considered alongside uncertainty, physical consistency and validation under relevant operating conditions.

8.2 Sources of Uncertainty

Uncertainty in materials degradation forecasting arises from several distinct sources. Aleatoric uncertainty reflects variability that is inherent to material behaviour and operating conditions, such as differences in microstructure, loading or environmental exposure. Epistemic uncertainty results from incomplete knowledge and limited training data and may therefore be reduced when additional representative information becomes available. Measurement uncertainty can arise from sensor noise, calibration errors and limitations in experimental techniques.

Further uncertainty can originate from model assumptions, incomplete descriptions of degradation mechanisms and changes in environmental or loading conditions. Dataset uncertainty is also important when training data do not adequately represent the materials or service environments encountered during deployment. These sources may occur simultaneously, making uncertainty an important property of the forecasting problem rather than simply an error associated with the AI algorithm.

8.3 Uncertainty Quantification

Uncertainty quantification provides a means of expressing how reliable an AI forecast is under given conditions. Bayesian methods, including approximations such as Monte Carlo dropout (Gal & Ghahramani, 2016), can represent uncertainty in model parameters and predictions, while ensemble approaches such as deep ensembles (Lakshminarayanan et al., 2017) assess variation across multiple independently trained models. Probabilistic forecasting can provide distributions of possible degradation outcomes rather than a single predicted value, allowing the likelihood of different future states to be considered.

Prediction intervals provide another practical way of communicating the expected range of future observations. Conformal approaches can also be used to construct prediction intervals with specified coverage under appropriate assumptions, offering distribution-free guarantees that do not rely on strong assumptions about the underlying model (Angelopoulos & Bates, 2021; Bates et al., 2021). These methods are particularly relevant to long-term degradation forecasting because uncertainty generally increases as predictions extend further from the observed data. Reporting this uncertainty can therefore provide a more realistic representation of forecast reliability than reporting a single point prediction, building on the broader tradition of placing probability distributions over network weights and predictions established by early Bayesian neural network formulations (MacKay, 1992).

8.4 Explainability

Explainability is important because materials engineers need to understand the basis of AI predictions, particularly when forecasts indicate increasing degradation or failure risk. Feature importance methods can identify which environmental, mechanical or material variables

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723149

contribute strongly to a prediction. Shapley-value-based methods such as SHAP provide a theoretically grounded way of attributing a prediction to individual input features (Lundberg & Lee, 2017). For imaging and sequential data, gradient-based saliency methods such as Grad-CAM can indicate which regions or portions of an input receive greater importance from the model (Selvaraju et al., 2017).

Attention mechanisms can also provide information about relationships that a transformer considers important, although attention should not automatically be interpreted as a causal explanation. Physics-based interpretability offers a complementary approach by examining whether model behaviour is consistent with established degradation mechanisms. Explainability is therefore most useful when it helps relate model predictions to measurable material behaviour rather than simply providing a visual representation of model activity.

8.5 From Explainability to Actionable Trust

Explainability alone does not establish that an AI forecast is trustworthy. A model may provide apparently understandable explanations while still producing inaccurate or physically inconsistent predictions. Trustworthy failure forecasting requires a combination of predictive accuracy, physical consistency, calibrated uncertainty and independent validation. Calibration is particularly important because reported confidence should correspond reasonably well to the actual frequency of prediction errors.

External validation is also necessary to determine whether a model remains reliable when applied to materials, environmental conditions or loading histories that were not used during model development. For practical applications, the most valuable forecasting systems will therefore be those that not only predict degradation accurately, but also indicate when their predictions are uncertain and provide evidence that their behaviour remains consistent with established materials science.

IX. TRANSLATIONAL CHALLENGES: FROM LABORATORY TO FIELD

9.1 The Long-Term Data Problem

Reliable long-term forecasting of materials degradation is limited by the availability of suitable longitudinal datasets. True failure events are relatively uncommon, while experiments designed to follow materials until failure can require substantial time, resources and controlled monitoring. Many studies therefore rely on shorter experimental periods or accelerated tests that capture only part of the degradation process. Censored observations, in which testing ends before failure occurs, further limit the information available for modelling time to failure (Cox, 1972).

These limitations make it difficult for AI models to learn the complete progression from early degradation to failure. A model may perform well when predicting within the period represented by the training data but have limited evidence for forecasting behaviour over much longer service periods. Developing sufficiently long and diverse degradation datasets is therefore essential for evaluating whether AI systems can genuinely predict failure rather than reproduce short-term patterns (Hamada et al., 2025).

9.2 Accelerated Aging vs Real-World Aging

Accelerated aging is widely used to obtain degradation data within practical experimental timescales. However, increasing temperature, moisture, chemical exposure or mechanical loading may alter the relative importance of degradation mechanisms compared with those operating under normal service conditions. Real materials are also exposed to variable and interacting stressors rather than constant laboratory conditions.

Consequently, performance under accelerated aging cannot automatically be assumed to represent real-world degradation. An AI model trained predominantly on accelerated experiments may learn relationships specific to those artificial conditions and fail to reproduce the variability and coupled stressor interactions encountered in service; this gap has been highlighted specifically for polymer systems, where classical accelerated-aging extrapolation methods such as the Arrhenius relationship and time–temperature superposition can break down under multi-mechanism coupling that machine learning approaches are only beginning to address (Zeng et al., 2025). Validation against realistic exposure histories is therefore necessary before such models can be considered reliable for field forecasting.

9.3 Cross-Material Generalization

Materials differ substantially in composition, microstructure, mechanical behaviour and dominant degradation mechanisms. A forecasting model developed for one material class may therefore not transfer directly to another. Alloys, polymers, concrete, composites and protective coatings can respond differently to temperature, moisture, chemical exposure and mechanical loading, even when subjected to similar external conditions.

Cross-material generalization requires models to distinguish relationships that are broadly applicable from those that are specific to a particular material system. Without sufficient representation of different material classes, a model may perform well within its training domain but show substantial performance loss when applied to unfamiliar materials. Transfer learning, in which knowledge learned in a source domain or task is reused to improve learning in a related but distinct target domain (Pan & Yang, 2010), and physics-informed approaches (Karpatne et al., 2017) may help address this problem, but their effectiveness must be demonstrated through independent validation rather than assumed from performance within the original dataset.

9.4 Cross-Environment Generalization

Materials experience substantially different environmental conditions across climates, industrial settings and operating regimes. Temperature and humidity cycles, chemical exposure, loading frequency and magnitude can vary considerably between service environments. These differences can change both the rate and mechanism of degradation.
An AI model trained under a narrow environmental range may therefore struggle when applied to conditions outside that range. This is especially important when the relationship between exposure and degradation is nonlinear. Reliable forecasting requires datasets that capture meaningful environmental variability and evaluation under conditions that differ from those used for model training. Cross-environment testing is consequently an important measure of whether an AI system has learned general degradation behaviour or simply adapted to a particular operating regime.

9.5 Dataset Bias and Representation

Dataset composition can strongly influence the behaviour of AI degradation models. Some materials, degradation mechanisms or failure modes may be heavily represented, while rare but important conditions receive little attention. Extreme environmental exposures and unusual failure pathways may also be underrepresented because they are difficult to reproduce or occur infrequently in service.

Sensor availability can introduce another form of bias. Materials or structures equipped with extensive monitoring systems may generate much richer datasets than those with limited sensing. As a result, models may perform better for well-monitored systems and poorly for conditions where observations are sparse. Addressing dataset bias therefore requires attention to the diversity of materials, degradation states, environments and measurement conditions represented during model development.

9.6 Laboratory-to-Field Translation

The transition from laboratory validation to reliable field forecasting remains a major unresolved challenge. Laboratory studies provide controlled conditions that are valuable for isolating degradation mechanisms, but field environments introduce variability in exposure, loading, material history, maintenance and operational conditions that may not be reproduced experimentally.

Consequently, strong laboratory performance does not establish that an AI model will remain accurate during real service. External validation using realistic operating conditions and independent datasets is necessary to determine whether the model can generalize beyond the environment in which it was developed. For materials degradation forecasting, successful laboratory-to-field translation therefore requires not only accurate models, but also representative data, appropriate uncertainty estimates and evidence of robust performance under changing real-world conditions. The translational challenges discussed above highlight the gap between promising AI performance under controlled research conditions and reliable degradation forecasting in real service environments. These challenges arise from limitations in data availability, model generalization, environmental representation, uncertainty and validation. Table 3 summarizes the major barriers to practical implementation, their potential consequences for AI-based materials degradation forecasting, and possible approaches for addressing them.

Table 3. Major Translational Barriers, Consequences and Proposed Solutions

Major barrier	Potential consequence	Proposed solution
Limited long-term degradation data	Poor prediction of long-term degradation and failure	Develop longitudinal datasets and standardized long-duration experiments
Accelerated-aging dependence	Laboratory predictions may not represent real service behaviour	Validate models using realistic and variable exposure histories
Limited cross-material generalization	Reduced performance on unfamiliar materials	Use diverse datasets, transfer learning and physics-informed approaches
Limited cross-environment generalization	Poor performance under unseen environmental conditions	Include broad environmental variability and independent external testing
Dataset bias and imbalance	Underprediction of rare degradation mechanisms and failure modes	Improve dataset diversity and representation of relevant conditions
Laboratory-to-field gap	Strong laboratory performance but unreliable field predictions	Conduct independent field validation under realistic operating conditions
Measurement uncertainty and missing data	Unstable state estimation and forecasting errors	Improve sensing, data-quality control and missing-data handling
Limited model interpretability and uncertainty assessment	Reduced confidence in predictions for engineering decisions	Integrate explainability, uncertainty quantification and calibration

X. AN INTEGRATED FRAMEWORK FOR MULTIMODAL DEGRADATION FORECASTING

An integrated framework for multimodal degradation forecasting should represent the progression from environmental exposure to material deterioration, prediction and engineering decision making. The first layer is the Environmental Exposure Layer, which captures temperature, moisture, chemical exposure, mechanical loading and the history of these conditions. This information provides the external context in which degradation occurs. The Material State Evolution Layer then represents the resulting changes within the material, including microstructural changes, chemical degradation, crack initiation, damage accumulation and deterioration of mechanical or functional properties. Linking exposure to evolving material state is essential because the same external condition can produce different outcomes depending on material history and prior damage.

The Multimodal Observation Layer provides measurements of this evolving state through sensors, imaging, mechanical testing, environmental monitoring and simulation outputs (Baltrušaitis et al., 2019). These observations are then supplied to the AI Forecasting Layer, which estimates future material condition, degradation trajectories, remaining useful life and failure probability. Importantly, the framework does not treat the AI forecast as an independent final output. A Physics and Uncertainty Layer constrains predictions using relevant physical knowledge (Raissi et al., 2019; Karniadakis et al., 2021) while quantifying uncertainty, assessing calibration and supporting interpretation of model behaviour (Lakshminarayanan et al., 2017; Angelopoulos & Bates, 2021). This provides a basis for determining whether a predicted degradation trajectory is physically plausible and sufficiently reliable for use.

The final Decision Layer translates these forecasts into practical engineering actions, including risk ranking, inspection prioritization, maintenance planning and early intervention. The framework therefore establishes a continuous pathway from environmental exposure through material state evolution and multimodal observation to AI-based forecasting and uncertainty assessment, ultimately supporting proactive management of degradation and failure risk. The proposed conceptual sequence is: environmental exposure → material state → multimodal sensing → AI forecast → uncertainty → failure probability → proactive intervention.

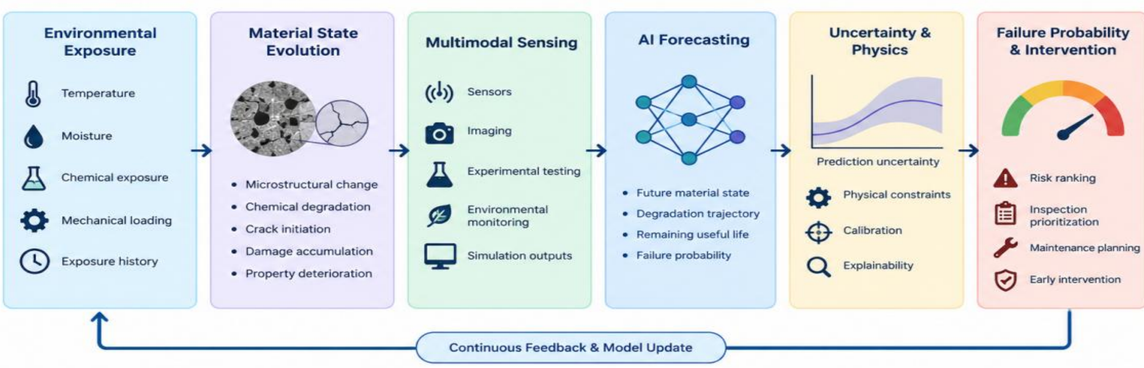

Figure 1. Integrated framework for multimodal degradation forecasting. The framework links environmental exposure and material state evolution with multimodal observations, AI-based forecasting, uncertainty assessment and physics-based constraints. Predicted degradation trajectories and failure probability support proactive inspection, maintenance and intervention, while continuous feedback allows the forecasting system to be updated as new material and environmental data become available.

XI. FUTURE DIRECTIONS AND RESEARCH PRIORITIES

Future research in AI-based materials degradation forecasting should place greater emphasis on the development of long-term and openly accessible benchmark datasets. Standardized experimental protocols and reporting practices would improve the comparability of models and make it easier to determine whether reported advances represent genuine improvements in forecasting capability. This is particularly important for failure prediction, where labelled failure data are often scarce. Self-supervised and few-shot learning could help extract useful representations from large amounts of unlabelled degradation data while reducing dependence on extensive labelled datasets.

Another important direction is the development of foundation models for materials degradation that can learn from diverse materials, environments and data modalities. Such models could provide a broader basis for transfer across material systems and operating conditions. Physics-constrained multimodal AI may further improve generalization by combining physical knowledge with information from sensors, imaging, experiments and simulations (Karniadakis et al., 2021). Continual and online learning will also become increasingly important as materials age and environmental conditions change, allowing forecasting systems to update their predictions as new observations become available.

Future studies should also adopt more consistent validation standards that include external validation, uncertainty quantification, calibration and cross-domain testing. Performance should be assessed not only on familiar datasets but also across different materials, environments and degradation conditions. Ultimately, these developments could move materials management toward increasingly autonomous predictive maintenance, following a progression from monitor → detect → predict → recommend → adapt. Such systems would shift the emphasis from responding to observed damage toward continuously anticipating degradation and supporting timely intervention.

XII. CONCLUSION

Materials failure under realistic service conditions is a dynamic and complex process driven by interacting environmental and mechanical stressors. Temperature, moisture, chemical exposure and loading can produce changes across multiple material scales, from microstructural and chemical alterations to visible damage and loss of functional performance. Consequently, reliable degradation assessment requires more than identifying damage after it has occurred.

AI offers an opportunity to move materials assessment from periodic damage detection toward continuous forecasting of degradation trajectories, remaining useful life and failure probability. However, achieving this transition requires more than increasingly complex model architectures. Reliable forecasting depends on high-quality longitudinal data, appropriate integration of complementary data modalities, incorporation of relevant physical constraints and explicit treatment of predictive uncertainty. Rigorous external validation across different materials, environments and operating conditions is equally important for establishing whether models can generalize beyond their development datasets.

The ultimate objective should therefore not be higher predictive accuracy alone. The value of AI for materials degradation lies in its ability to provide sufficiently early, physically credible and trustworthy forecasts that can support safer, more efficient and proactive intervention throughout the service life of materials and structures.

Declarations

Funding
No specific funding was received for this work.

Conflict of Interest
The authors declare that they have no conflict of interest.

Ethics Approval
Not applicable. This manuscript is a critical review of previously published literature and does not involve human participants, animals, or primary experimental data.

Consent to Participate
Not applicable.

Consent for Publication
Not applicable.

Data Availability
No new datasets were generated or analyzed in this study. All information discussed is derived from previously published literature.

Author Contributions
The authors contributed to the conception and development of the review, literature analysis, interpretation of the evidence, and preparation of the manuscript. All authors contributed to critical revision of the manuscript, approved the final version, and agreed to be accountable for the content of the work.

REFERENCES

[1] Angelopoulos, A. N., & Bates, S. (2021). A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *arXiv:2107.07511*. arXiv

[2] Atha, D. J., & Jahanshahi, M. R. (2018). Evaluation of deep learning approaches based on convolutional neural networks for corrosion detection. *Structural Health Monitoring, 17*(5), 1110–1128. SAGE

[3] Baltrušaitis, T., Ahuja, C., & Morency, L.-P. (2019). Multimodal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 41*(2), 423–443. IEEE

[4] Bates, S., Angelopoulos, A., Lei, L., Malik, J., & Jordan, M. I. (2021). Distribution-free, risk-controlling prediction sets. *Journal of the ACM, 68*(6), 43:1–43:34. ACM

[5] Breiman, L. (2001). Random forests. *Machine Learning, 45*(1), 5–32. Springer

[6] Cha, Y.-J., Choi, W., & Büyüköztürk, O. (2017). Deep learning-based crack damage detection using convolutional neural networks. *Computer-Aided Civil and Infrastructure Engineering, 32*(5), 361–378. Wiley

[7] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794). ACM. ACM

[8] Cox, D. R. (1972). Regression models and life-tables. *Journal of the Royal Statistical Society: Series B, 34*(2), 187–220.

[9] Fuhaid, A. F. A., & Niaz, A. (2022). Carbonation and corrosion problems in reinforced concrete structures. *Buildings, 12*(5), 586. MDPI

[10] Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of the 33rd International Conference on Machine Learning* (pp. 1050–1059). PMLR

[11] Grieves, M., & Vickers, J. (2017). Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In F.-J. Kahlen, S. Flumerfelt, & A. Alves (Eds.), *Transdisciplinary Perspectives on Complex Systems* (pp. 85–113). Springer. Springer

[12] Hamada, A., Elyamny, S., Abd-Elaziem, W., Elkatatny, S., Darwish, M. A., Sebaey, T. A., Järvenpää, A., Vineesh, K. P., & Elsheikh, A. H. (2025). Advancing fatigue life prediction with machine learning: A review. *Materials Today Communications, 45*, 111525.

[13] Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *Science, 313*(5786), 504–507. Science

[14] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation, 9*(8), 1735–1780. MIT Press

[15] Jia, H., Qiao, G., & Han, P. (2022). Machine learning algorithms in the environmental corrosion evaluation of reinforced concrete structures — a review. *Cement and Concrete Composites, 133*, 104725. ScienceDirect

[16] Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. *Nature Reviews Physics, 3*(6), 422–440. Nature

[17] Karpatne, A., Atluri, G., Faghmous, J. H., Steinbach, M., Banerjee, A., Ganguly, A., Shekhar, S., Samatova, N., & Kumar, V. (2017). Theory-guided data science: A new paradigm for scientific discovery from data. *IEEE Transactions on Knowledge and Data Engineering, 29*(10), 2318–2331.

[18] Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems, 25*, 1097–1105.

[19] Lakshminarayanan, B., Pritzel, A., & Blundell, C. (2017). Simple and scalable predictive

uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems 30* (pp. 6402–6413).

[20] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature, 521*(7553), 436–444. Nature

[21] Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems 30* (pp. 4765–4774).

[22] MacKay, D. J. C. (1992). A practical Bayesian framework for backpropagation networks. *Neural Computation, 4*(3), 448–472. MIT Press

[23] Nash, W., Zheng, L., & Birbilis, N. (2022). Deep learning corrosion detection with confidence. *npj Materials Degradation, 6*, 26. Nature

[24] Osinubi, O. O., Njemanze, E. C., Dudzilah, G., Enobakhare, B. O., & Adeoba, M. I. (2026). A review of health-aware modeling and control strategies for battery energy storage systems. *Computational Intelligence in Modern Power Systems, 1*(1), 24–34. Crossref

[25] Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering, 22*(10), 1345–1359. IEEE

[26] Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics, 378*, 686–707. ScienceDirect

[27] Rasmussen, C. E., & Williams, C. K. I. (2006). *Gaussian processes for machine learning*. MIT Press.

[28] Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE International Conference on Computer Vision* (pp. 618–626).

[29] Tuegel, E. J., Ingraffea, A. R., Eason, T. G., & Spottswood, S. M. (2011). Reengineering aircraft structural life prediction using a digital twin. *International Journal of Aerospace Engineering, 2011*, 154798. Wiley

[30] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems 30* (pp. 5998–6008).

[31] Wang, Q., Huang, B., Gao, Y., & Jiao, C. (2025). Current status and prospects of digital twin approaches in structural health monitoring. *Buildings, 15*(7), 1021. MDPI

[32] Zeng, B., Wu, S., & Yao, S. (2025). Machine learning-driven paradigm for polymer aging lifetime prediction: Integrating multi-mechanism coupling and cross-scale modeling. *Polymers, 17*(22), 2991. MDPI