

Review

Scientific Machine Learning for Coating Degradation and Corrosion Prediction: A Critical Review and Thermodynamic Research Perspective

Luis Rojas-Valdivia * , Paola Moraga, Alvaro Peña  and José García * 

Escuela de Ingeniería de Construcción y Transporte, Pontificia Universidad Católica de Valparaíso, Valparaíso 2362804, Chile; paola.moraga@pucv.cl (P.M.); alvaro.pena@pucv.cl (A.P.)

* Correspondence: luis.rojas.v@mail.pucv.cl (L.R.-V.); jose.garcia@pucv.cl (J.G.)

Abstract

Coating degradation and corrosion prediction are moving from empirical regression toward scientific machine learning (SciML), which embeds governing equations, geometry, and uncertainty into data-driven models. This critical review examines physics-informed neural networks (PINNs), neural operators, phase-field hybrids, graph neural networks, and probabilistic models relevant to protective coatings and localized corrosion. The synthesis identifies five persistent limitations: dimensionally inconsistent composite losses; incomplete enforcement of conservation and thermodynamic irreversibility; inadequate treatment of anomalous transport and material memory; loss of topological information during pit nucleation and coalescence; and weak out-of-distribution validation. The evidence is organized by governing mechanism, physical incorporation, geometric representation, uncertainty treatment, and validation design, and it is consolidated in a study-level comparison that distinguishes experimentally demonstrated capabilities from simulation-validated and conceptual ones across materials, coating classes, and corrosive environments. On this basis, this paper proposes a Thermo-Fractional Graph Phase-Field Neural Operator (TFG-PFNO) as a research framework that combines fractional memory, graph-based geometry, phase-field interface evolution, metriplectic dynamics, and calibrated uncertainty. The framework is a testable hypothesis rather than a validated model. Its principal implication is that credible corrosion digital twins must preserve physical admissibility and propagate uncertainty to structural risk, inspection planning, and maintenance decisions, rather than optimize interpolation accuracy alone.

Keywords: scientific machine learning; physics-informed neural networks; protective coatings; corrosion; phase-field modeling; neural operators; graph neural networks; uncertainty quantification; digital twins



Academic Editor: Yongxian Huang

Received: 16 July 2026

Revised: 27 July 2026

Accepted: 28 July 2026

Published: 3 August 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Protective coatings delay corrosion by controlling the transport of water, oxygen, and aggressive ions; stabilizing interfacial chemistry; and redistributing mechanical stresses. They remain the most widely deployed corrosion-mitigation technology in atmospheric, marine, and buried service, and their barrier, inhibitive, and sacrificial protection mechanisms have been consolidated in classical accounts of anticorrosive coating science [1–3]. Their degradation is therefore not a purely surface phenomenon and cannot be represented adequately by a single scalar such as mass loss, average corrosion rate, or maximum pit

depth. The measured response emerges from a coupled, open, nonequilibrium system in which ionic transport, charge transfer, water uptake, polymer relaxation, residual stress, interfacial decohesion, localized dissolution, and fracture evolve over widely separated length and time scales. Studies of organic coatings illustrate this coupling: physically based equivalent-circuit analysis, neural interpretation of EIS, and unsupervised grouping of impedance spectra all depend on the measurement model and degradation state [4–7]. Recent corrosion-oriented machine learning reviews confirm the rapid growth of data-driven prediction, but they also show that model performance is difficult to compare across materials, environments, sensing modalities, and validation protocols [8–12].

A useful mathematical description treats experimentally accessible quantities—potentials, impedance spectra, thickness maps, surface images, and mass loss—as partial observations of a distributed latent state that gathers species concentrations, electrochemical potential, coating integrity, displacement, damage, and temperature. To keep this introduction focused on motivation and significance, the corresponding state-space formulation, its governing operator equation, and the baseline physics-informed learning objective are consolidated in Section 2, Equations (1)–(3) and Figure 1. The consequence matters more than the formalism: the learning task changes from tabular regression to the inference of fields, operators, and latent parameters that should remain causal, conservative, objective, and dissipative.




Figure 1. Baseline reaction–diffusion PINN for corrosion modeling. Space–time coordinates, material parameters, and environmental variables are mapped to electrochemical fields; automatic differentiation generates reaction, transport, and conservation residuals. The diagram is conceptual and does not report experimental performance.

Conventional multiphysics models provide interpretability but become expensive to calibrate when observations are sparse, indirect, and heterogeneous. Industrial data often combine electrochemical potentials, electrochemical impedance spectroscopy (EIS), ultrasonic thickness measurements, surface images, environmental histories, and pointwise mass-loss measurements. Conversely, a purely data-driven model can fit correlations in these measurements without preserving mass balance, charge balance, positivity, or the monotonic evolution of irreversible variables. Random train–test splits can further exaggerate performance when measurements from the same specimen, asset, coating batch, or exposure campaign appear in both sets [13–15].

Neural differential equation solvers predate the modern PINN formulation [16,17]. Scientific machine learning (SciML) extends this line of work by incorporating differential equations, variational principles, symmetries, and constitutive knowledge into the learning process [18–21]. In a physics-informed neural network (PINN), the unknown state is approximated by a differentiable surrogate $\hat{\mathbf{z}}_\theta(\mathbf{x}, t, \mu)$, automatic differentiation supplies the derivatives required by the governing residuals, and data, equation, boundary, and initial terms are combined in a single training objective [22,23], stated formally in Equation (3) of Section 2. This approach is attractive because data and physics can be combined in one differentiable model, but a low aggregate loss does not establish physical admissibility.

PINNs can converge unevenly across loss terms, particularly for stiff systems, sharp interfaces, multiscale transients, and poorly scaled parameters [24–27]. This limitation is especially consequential for localized corrosion. A model can reproduce an average field while smoothing the maximum pit depth, eliminating a narrow transport path, or misrepresenting pit coalescence. Gradient-enhanced residuals, temporal decomposition, adaptive loss balancing, and adaptive sampling can improve optimization [28–33], but they do not replace conservation checks, thermodynamic constraints, topology-sensitive metrics, or independent validation.

The central question is thus not whether a neural network can interpolate corrosion data, but what mathematical structure is required for a learned model to represent an irreversible, multiscale, and partially observed material system. Five issues are particularly important: (i) nonlocal memory associated with trapping and anomalous diffusion; (ii) coupled electrochemical, transport, phase-field, and mechanical dynamics; (iii) evolving geometry and topology during defect nucleation and coalescence; (iv) epistemic uncertainty caused by non-identifiable parameters and limited data; and (v) thermodynamic irreversibility. Each issue changes both the admissible architecture and the meaning of model error.

The practical stakes justify a dedicated critical review. The direct global cost of corrosion has been estimated at US\$2.5 trillion per year, close to 3.4% of the global gross domestic product, and protective coatings absorb the largest share of the corresponding mitigation effort [1,34]. Earlier surveys examined deep learning for materials degradation in general [35], corrosion machine learning from a data-oriented standpoint [8], or specific domains such as marine assets, electrochemical identification, and smart nanocomposite coatings [9,11,12]. None of these contributions, however, classifies the literature by the computational mechanism through which physical knowledge enters the model, nor does any separate systematically the capabilities that are demonstrated experimentally from those validated only against numerical simulations or still conceptual. Closing that methodological gap for coating degradation and corrosion prediction is the specific purpose, and in our view the principal significance, of the present review.

This article therefore presents a critical, mechanism-centered review of PINNs, neural operators, graph neural networks (GNNs), phase-field hybrids, and probabilistic SciML for coating degradation and corrosion prediction. It makes five contributions. First, it organizes the literature by how physical knowledge is incorporated rather than by model name. Second, it defines a minimum physics and validation stack for corrosion applications. Third, it identifies which claims are supported by verifiable corrosion-specific evidence and which remain extrapolations from general SciML. Fourth, it proposes the Thermo-Fractional Graph Phase-Field Neural Operator (TFG–PFNO) as a testable research framework. Fifth, it provides a study-level evidence map presented in Section 3.6 that records, for each corrosion- or coating-specific SciML study, the material system, exposure environment, degradation mechanism, input data, architecture, physical constraints, validation strategy, principal finding, and limitation, so that experimentally demonstrated capabilities can be distinguished from simulation-based or purely conceptual ones. The proposed architecture is not presented as a benchmarked model; it is a structured hypothesis intended to guide future implementation and validation.

The remainder of the article is organized as follows: Section 2 defines the review scope, the literature search and screening procedure, the analytical framework, and the baseline mathematical formulation. Section 3 synthesizes the evidence by governing mechanism, physical incorporation, geometric representation, uncertainty treatment, and validation design; it then adds the study-level comparison and an explicit discussion of coating classes, exposure environments, measurement techniques, datasets, and industrial barriers.

Section 4 develops the TFG–PFNO research framework and its falsification plan. Section 5 distills the evidence-based research gaps and future directions.

2. Review Scope and Evidence Quality Procedure

2.1. Review Design

This work is a critical review and methodological perspective, not a systematic review or meta-analysis. The distinction is deliberate: the relevant studies span field regression, image analysis, inverse parameter identification, partial differential equation (PDE) solvers, phase-field simulations, structural capacity models, and probabilistic prognosis. Their outcomes and validation designs are too heterogeneous for a defensible pooled effect size. The synthesis therefore focuses on verifiable physical and computational claims: which equations or constraints were implemented, how they entered the model, what data supported the model, and whether evaluation extended beyond interpolation.

The scope was organized around three search blocks: (i) SciML methods (“physics-informed”, PINN, neural operator, physics-guided learning, and hybrid modeling); (ii) coating and corrosion processes (protective coatings, pitting, electrochemistry, stress corrosion cracking, corrosion fatigue, and atmospheric or marine corrosion); and (iii) mathematical mechanisms (reaction–diffusion, phase field, fractional transport, graph learning, inverse problems, and uncertainty quantification). The main temporal emphasis was 2019–July 2026, with earlier publications retained when they established indispensable mathematical or computational foundations.

Table 1 states the decision framework used for inclusion and interpretation. General SciML papers were retained only when they contributed a mechanism directly transferable to coating degradation, such as exact boundary enforcement, adaptive sampling, domain decomposition, operator learning, energy-based training, conservative message passing, or Bayesian inference. Corrosion papers based only on image classification were used selectively and were not treated as physics-informed unless the method linked observations to a physical state, temporal evolution, or a governing constraint.

Table 1. Framework used to select and interpret evidence in this critical review.

Dimension	Included Evidence	Interpretive Rule
Physical domain	Coating transport, electrochemistry, pitting, interface evolution, damage, structural capacity, and corrosion monitoring	A corrosion label alone was insufficient; the modeled state and mechanism had to be identifiable.
Physics mechanism	PDE residuals, conservation, hard constraints, variational energies, constitutive laws, symmetry, and thermodynamic inequalities	Classification was based on the implemented mechanism, not on an author’s use of the term “physics-informed”.
Data and validation	Experiments, field sensing, independent simulations, analytical solutions, and documented train–validation–test partitions	Evidence from a new specimen, environment, geometry, or time horizon received greater weight than random record-level splits.
Computational reporting	Architecture, optimizer, collocation or mesh strategy, baselines, error metrics, code, and data availability	Missing information was treated as a limitation rather than inferred from common practice.
Claim strength	Consolidated, moderate, emerging, or conceptual	Corrosion-specific experimental support was distinguished from transfer arguments based on general SciML literature.

2.2. Search Strategy, Screening, and Corpus Composition

Although this work is a critical review rather than a systematic review, the literature selection process is reported here so that it can be reproduced to a reasonable approximation. Searches were performed in Scopus, the Web of Science Core Collection, ScienceDirect, IEEE

Xplore, and SpringerLink, complemented by Google Scholar, arXiv, and SSRN for preprints and conference material. The search window extended from January 2015 to July 2026, with the primary emphasis on 2019–2026; earlier sources were admitted only when they establish indispensable mathematical, electrochemical, or computational foundations. Representative Boolean strings combined the three thematic blocks defined above, for example: (“physics-informed” OR “PINN” OR “neural operator” OR “physics-guided” OR “graph neural network”) AND (coating* OR corrosion OR pitting OR delamination OR electrochem*) AND (diffusion OR “phase field” OR fractional OR uncertainty OR “digital twin”).

Screening proceeded in two stages. Titles and abstracts were first examined against four inclusion criteria: (i) a peer-reviewed publication or publicly archived preprint with an identifiable methodology; (ii) an identifiable physical state, mechanism, or governing constraint (a corrosion label alone was insufficient); (iii) verifiable bibliographic metadata, normally a resolvable DOI; and (iv) availability in English. Exclusion criteria were: purely cosmetic uses of the term “physics-informed”; image-only classification with no link to a physical state or temporal evolution (retained only selectively as monitoring evidence); inaccessible or inconsistent metadata; and duplicated content across venues. In approximate figures, the searches returned somewhat more than five hundred unique records after deduplication; about one hundred ninety documents were assessed in full text; and 124 sources satisfied all criteria and are cited in this review, of which roughly forty percent address corrosion or coating systems directly while the remainder provide mathematical, computational, or statistical foundations. These counts are approximate by necessity, because the search was iterative and complemented by backward and forward snowballing from key papers; bibliographic records were audited individually, including replacements and exclusions.

2.3. Source Screening

Sources were retained when their bibliographic information was internally consistent and their content was relevant to the claim under discussion. Publisher records and DOI landing pages were consulted when metadata were ambiguous, and entries that could not be resolved consistently were excluded. This screening was intended to reduce citation and metadata errors; it was not a substitute for evaluating study design, physical assumptions, data provenance, or validation quality.

Beyond the approximate screening counts reported above, the review avoids bibliometric growth rate claims and pooled effect sizes that cannot be reconstructed independently. It also avoids treating conceptual figures or synthetic curves as empirical results.

2.4. Analytical Framework

The literature was interpreted along six axes: governing physics, mechanism of physics incorporation, geometric representation, optimization strategy, uncertainty treatment, and validation design. A claim was considered physically substantive only when it connected (i) a governing equation, constitutive law, variational principle, or inequality; (ii) an identifiable computational mechanism; (iii) data or simulations with a stated origin; and (iv) a performance or admissibility test.

The synthesis distinguishes four levels of incorporation. *Physics-guided* models use physical knowledge to select variables or generate data. *Physics-informed* models place equations or constraints in the objective function. *Physics-encoded* models enforce a property through the architecture, for example antisymmetric fluxes, equivariance, positivity, or exact boundary satisfaction. Finally, a model is described here as *thermodynamically compiled* when its parameterization makes an inadmissible sign of dissipation structurally impossible

within the stated assumptions. This final category is an analytical proposal of the present review, not an established community standard.

Risk of overstatement was assessed qualitatively in eight domains: limited or non-representative data; non-independent sampling; information leakage; circular validation against the data-generating simulator; weak baselines; lack of out-of-distribution (OOD) tests; uncalibrated uncertainty; and incomplete reproducibility. Conclusions supported by multiple independent studies and corrosion-specific validation are described more strongly than conclusions inferred from mathematical compatibility alone.

2.5. Baseline Mathematical Formulation

Following the reorganization of this article, the state-space formulation and the baseline physics-informed objective that anchor the entire synthesis are stated here rather than in the introduction. A useful mathematical description treats experimentally accessible quantities as partial observations of a distributed latent state,

$$\mathbf{z}(\mathbf{x}, t) = [c_1, \dots, c_{N_s}, \psi, \phi, \mathbf{u}, d, \vartheta]^\top, \quad (1)$$

where c_i denotes the concentration of chemical species i , ψ is the electrochemical potential, ϕ represents coating or interface integrity, $\mathbf{u}$ is displacement, d is an internal damage variable, and ϑ is temperature. In abstract form, the state evolves according to

$$\mathcal{M}(\mathbf{z}, \boldsymbol{\mu}) \partial_t \mathbf{z} = \mathcal{A}(\mathbf{z}, \nabla \mathbf{z}, \nabla^2 \mathbf{z}; \boldsymbol{\mu}) + \mathcal{S}(\mathbf{x}, t; \boldsymbol{\mu}), \quad (2)$$

subject to initial conditions, boundary conditions, conservation laws, and constitutive inequalities. The parameter vector $\boldsymbol{\mu}$ may include diffusivities, permeabilities, conductivities, reaction constants, activation energies, coating properties, and fracture parameters.

In a PINN, the unknown state is approximated by $\hat{\mathbf{z}}_\theta(\mathbf{x}, t, \boldsymbol{\mu})$, and automatic differentiation supplies the derivatives required by the governing residuals [22,23]. A conventional objective is

$$\mathcal{L}_{\text{PINN}} = \lambda_{\text{data}} \mathcal{L}_{\text{data}} + \lambda_{\text{PDE}} \mathcal{L}_{\text{PDE}} + \lambda_{\text{BC}} \mathcal{L}_{\text{BC}} + \lambda_{\text{IC}} \mathcal{L}_{\text{IC}}. \quad (3)$$

Figure 1 summarizes this baseline for a reaction–diffusion corrosion problem. The capacity of such formulations to recover hidden fields from indirect observations was demonstrated prominently in fluid mechanics [36], and it motivates the inverse corrosion problems reviewed in Section 3.

3. Critical Synthesis

This section proceeds from governing physics to architectures, uncertainty, and validation; the individual studies discussed along the way are then consolidated in the study-level comparison of Section 3.6 and connected to coating practice in Section 3.7.

3.1. Minimum Governing Structure for Coating Degradation

The governing structure required by a SciML model depends on the decision it is intended to support. A short-horizon regression of corrosion potential may not require an explicit moving interface; prediction of pit growth, delamination, or remaining structural capacity usually does. At minimum, ionic transport should respect local species balance,

$$\partial_t c_i + \nabla \cdot \mathbf{J}_i = R_i(\mathbf{c}, \psi, \phi, \vartheta, d), \quad (4)$$

$$\mathbf{J}_i = -D_i(\phi, d, \vartheta) \nabla c_i - \frac{z_i F}{R \vartheta} D_i c_i \nabla \psi + c_i \mathbf{v}, \quad (5)$$

where R_i is the reaction source, D_i is an effective diffusivity, z_i is ionic valence, and $\mathbf{v}$ is an effective convective velocity. Charge consistency may be represented by Poisson's equation or, at the appropriate scale, by electroneutrality. Interfacial kinetics introduce further nonlinearities through Butler–Volmer-type relations. Recent corrosion-specific PINN studies have begun to use such physical structure for pit depth, atmospheric corrosion, and subsea pipeline prediction, but the level of closure varies substantially [37–40]. Related physics-guided studies use chloride transport equations for inverse diffusivity identification and oxide growth priors for liquid-metal corrosion prediction [41,42], demonstrating that the encoded physics can range from field equations to monotonic kinetic structure.

Coating transport may depart from the Fickian ideal because of tortuosity, trapping, polymer relaxation, microvoids, evolving saturation, and damage-assisted pathways. A compact representation is

$${}^C D_t^\alpha S = \nabla \cdot (\mathbf{D}_\beta \nabla^\beta S) + R_S, \quad 0 < \alpha \leq 1, \quad 1 < \beta \leq 2, \quad (6)$$

where the Caputo derivative ${}^C D_t^\alpha$ represents effective temporal memory and ∇^β represents a possible spatially nonlocal operator. Fractional dynamics provide a parsimonious description of distributed waiting times and anomalous transport [43–47]. Their use is meaningful only when the fractional order is identifiable and improves extrapolation relative to a classical model. Otherwise, α or β can become another flexible fitting parameter.

Localized dissolution and interface recession are naturally represented by a phase field $\phi \in [0, 1]$. A generic free energy is

$$\mathcal{F} = \int_{\Omega} \left[f_{\text{chem}}(\mathbf{c}, \psi, \phi) + \frac{\kappa_\phi}{2} |\nabla \phi|^2 + g(\phi) \Psi_{\text{el}}(\epsilon, d) + \frac{\kappa_d}{2} |\nabla d|^2 \right], d\Omega. \quad (7)$$

Phase-field models can represent pit nucleation, growth, and coalescence without explicit front tracking. The corrosion literature has progressed from activation- and diffusion-controlled pitting to galvanic coupling, quantitative kinetics, crystallographic anisotropy, electro-chemo-mechanical interaction, and temperature-dependent electrochemistry [48–54]. These computational advances rest on classical mechanistic foundations that any learned interface model must still reproduce: the critical factors governing pit initiation, stability, and repassivation [55] and the point defect description of passive film breakdown [56]. Dissolution-driven stress corrosion cracking has likewise been given a rigorous variational phase-field formulation [57], which provides a mechanistic benchmark against which future SciML couplings of corrosion and fracture should be tested. Coating-specific extensions are now beginning to connect phase-field simulations with accelerated or machine-learned prediction [58,59]. However, a learned evolution must also satisfy an admissibility condition. For an isothermal system,

$$\mathcal{D} = \mathcal{P}_{\text{ext}} - \frac{d\mathcal{F}}{dt} \geq 0, \quad (8)$$

where $\mathcal{P}_{\text{ext}}$ is external power. Penalizing violations of Equation (8) is useful, but encoding positive mobilities and dissipative operators in the architecture provides a stronger guarantee.

Figure 2 illustrates the separation of memory, reversible transport, and irreversible dissipation. Figure 3 shows the complementary use of staggered blocks for a tightly coupled phase-field corrosion problem. Both are conceptual design diagrams; they do not establish that a particular implementation outperforms a conventional solver.




Figure 2. Conceptual thermodynamic–fractional PINN. A memory operator augments the latent state, while reversible and dissipative blocks are separated before the governing and data constraints are evaluated.




Figure 3. Conceptual staggered phase-field PINN for coating or metal degradation. Alternating electrochemical and interface updates can reduce optimization interference in a strongly coupled system.

3.2. From Physics-Guided to Physics-Encoded Learning

The phrase “physics-informed” currently covers methods with markedly different guarantees. Table 2 separates four levels according to the computational mechanism that carries the physical knowledge. The distinction matters in corrosion because an environmental feature such as temperature or chloride concentration is physically meaningful, but its presence as an input does not enforce a balance law. Likewise, a PDE penalty reduces average residuals but does not ensure positivity, conservation after discretization, or nonnegative dissipation outside the collocation set.

Hard enforcement is not universally superior. Exact boundary parameterizations can become difficult on irregular geometries, and an incorrectly specified constitutive restriction can bias the solution more strongly than a soft penalty. The practical objective is to encode mature and non-negotiable properties, such as positivity or internal flux cancellation, while retaining residual or probabilistic treatment for uncertain closure relations. Methods for exact boundary imposition, domain decomposition, variational training, and energy-based mechanics provide useful components for this hierarchy [60–67], with the Deep Ritz method having established the variational energy route on which several of these formulations build [68]. Hybrid FEM–neural formulations provide a further route when a trusted discretization should remain explicit [69]. Conversely, sparse equation discovery and inverse learning of constitutive relationships can identify missing model terms, but their output must still be checked for identifiability and physical admissibility [70,71].

Table 2. Levels of physical incorporation used in this review.

Level	Mechanism	Typical Coating/Corrosion Use	Guarantee and Limitation
Physics-guided	Variable selection, feature engineering, synthetic data, or physically motivated augmentation	Environmental corrosion rate regression; image or sensor features linked to exposure conditions	Improves relevance but does not enforce governing physics.
Physics-informed	PDE, boundary, initial, energy, or inequality residuals in the loss	Inverse diffusivity, pit-depth evolution, reaction–diffusion fields, or phase-field emulation	Encourages compliance at sampled points; sensitive to scaling, optimization, and collocation coverage.
Physics-encoded	Exact constraints, conservative fluxes, equivariance, positivity transforms, or hard boundary parameterizations	Mass/charge-conservative graph updates, positive diffusivity, exact essential conditions	Provides a structural property within stated assumptions; may reduce flexibility or complicate training.
Compiled physics	Antisymmetric reversible operator plus symmetric positive semidefinite dissipative operator and degeneracy conditions	Coupled transport, interface evolution, and damage with irreversible material loss	Makes the sign of dissipation structural; remains a proposed standard that requires corrosion-specific validation.

3.3. Loss Scaling and Optimization Pathologies

A composite corrosion loss often combines quantities with different units and natural magnitudes. Adding an electrochemical residual in $\text{mol m}^{-3} \text{s}^{-1}$, a mechanical equilibrium residual in N m^{-3} , and a dimensionless data error without normalization gives numerical scale, rather than physical importance, control over optimization. A preferable objective is

$$\mathcal{L}(\theta) = \lambda_D \mathcal{L}_D + \sum_k \lambda_k \left\| \frac{\mathcal{R}_k}{s_k + \epsilon} \right\|_2^2 + \lambda_{\text{th}} \left\| \frac{\langle -\mathcal{D} \rangle_+}{s_D + \epsilon} \right\|_2^2 + \lambda_{\text{topo}} \mathcal{L}_{\text{topo}} + \lambda_{\text{cal}} \mathcal{L}_{\text{cal}}, \quad (9)$$

where s_k is a physically interpretable scale for residual k . The weights may be adapted using gradient norms, neural tangent kernel information, attention, or multi-objective criteria, but lower bounds should prevent a physically essential term from being silenced. Known PINN failure modes include gradient imbalance, spectral bias, causal inconsistency, and inadequate sampling near sharp features [25–28,30,31,72,73].

Optimization diagnostics should therefore report individual residual histories, not only total loss. They should also evaluate physical observables that are not direct training targets, including total mass or charge imbalance, free energy rate, maximum pit depth, lost volume, interface position, and structural capacity. Synthetic training curves can illustrate these diagnostics, but they should not be presented as evidence of model performance unless generated by a documented experiment or benchmark.

3.4. Neural Operators and Geometric Learning

A PINN trained for a single geometry and parameter set approximates one solution instance. Corrosion digital twins instead require a map from initial fields, environmental histories, material properties, and geometry to future degradation fields,

$$\mathcal{G} : [\mathbf{z}_0, \mathbf{s}_{0:t}, \boldsymbol{\mu}, \Omega] \longmapsto \mathbf{z}(\mathbf{x}, t). \quad (10)$$

DeepONet, Fourier neural operators (FNOs), and physics-informed neural operators learn such maps between function spaces [74–77]. They can amortize the cost of repeated evaluation across environmental histories or parameter fields. Their main limitation for localized corrosion is geometric: standard spectral operators are most convenient on regular grids, while pits, pores, cracks, coating defects, and inspection networks form irregular and evolving structures.

GNNs offer a complementary representation by assigning concentrations, potentials, phase values, material identifiers, or sensor states to nodes and geometric or constitutive information to edges [78–81]. Conservation can be encoded by antisymmetric edge fluxes,

$$\hat{\mathbf{J}}_{ij} = -\hat{\mathbf{J}}_{ji}, \quad c_i^{n+1} = c_i^n - \frac{\Delta t}{V_i} \sum_{j \in \mathcal{N}(i)} A_{ij} \hat{\mathbf{J}}_{ij}^n + \Delta t, R_i^n. \quad (11)$$

This construction makes internal transfer cancel exactly. Equivariance can additionally improve transfer between orientations or discretizations. Yet geometric learning alone does not supply irreversibility: Hamiltonian or symplectic networks describe reversible dynamics, whereas corrosion requires a dissipative channel [82,83]. The appropriate comparison is therefore not simply CNN versus GNN, but local grid-based models versus geometry-aware models equipped with the conservation and dissipation structure required by the application.

Table 3 summarizes the principal inductive biases and failure modes. No architecture is universally best; the scientific task, geometry, and validation regime should determine the choice.

Table 3. Comparison of SciML model families for coating degradation and corrosion.

Model Family	Useful Inductive Bias	Promising Use	Principal Risk
PINN	Continuous coordinates and automatic differentiation	Inverse parameters, sparse observations, and single-instance PDE solutions	Loss imbalance, stiffness, spectral bias, and expensive retraining.
Phase-field PINN	Diffuse interface and variational structure	Pit growth, interface recession, coating fracture, and coalescence	High-order derivatives and strongly coupled optimization.
Neural operator	Maps between input and solution functions	Many-query prediction across histories, parameters, or initial states	Grid dependence, data demand, and limited geometric transfer.
GNN or mesh network	Irregular geometry, connectivity, and local message passing	Pits, pores, cracks, interfaces, meshes, and sensor networks	Conservation and irreversibility are not automatic.
Hybrid multifidelity model	Combines simulation, laboratory, and field information	Digital twins and online state updating	Simulator bias, fidelity imbalance, and domain shift.
Bayesian or ensemble SciML	Predictive distributions and parameter uncertainty	Risk-informed prognosis and sensor planning	High cost and misleading intervals if calibration is not tested.


3.5. Uncertainty, Multifidelity Learning, and Application Evidence

Uncertainty in coating degradation has at least four sources: environmental variability, measurement noise, uncertain material parameters, and model discrepancy. These sources should not be collapsed into one empirical band, because deep learning and SciML uncertainty methods differ in what they represent and in how calibration should be assessed [84,85]. A probabilistic SciML model seeks a posterior over states and parameters,

$$p(\mathbf{z}, \boldsymbol{\mu} \mid \mathcal{D}, \mathcal{M}) \propto p(\mathcal{D} \mid \mathbf{z}, \boldsymbol{\mu}) p(\mathbf{z} \mid \boldsymbol{\mu}, \mathcal{M}) p(\boldsymbol{\mu}), \quad (12)$$

and propagates that posterior to a decision variable such as time to failure or failure pressure. Bayesian PINNs and related probabilistic methods provide a basis for noisy inverse problems [86]; emerging corrosion fatigue work demonstrates the direct relevance of adaptive Bayesian physics-informed models [87]. Scalable approximations such as Monte Carlo dropout and deep ensembles offer practical baselines for epistemic uncertainty, although their intervals must be calibration-tested before any operational use [88,89]. Multifidelity information fusion can then combine broad low-fidelity coverage with sparse high-fidelity observations, provided that cross-fidelity discrepancy is modeled rather than ignored [90,91]. When simulator and experiment disagree systematically, the discrepancy should be represented explicitly in the sense of Bayesian model calibration [92], rather than silently absorbed into network weights, because an absorbed bias reappears unpredictably outside the calibration domain.


Figure 4 shows how low-fidelity simulations and high-fidelity observations can enter a shared operator model. Figure 5 extends the idea to hierarchical uncertainty propagation and value-of-information-based sensor selection. These diagrams emphasize a workflow rather than a particular estimator.



Operator target: environmental history → full corrosion and coating-damage fields

Figure 4. Conceptual multifidelity neural operator digital twin. Phase-field or finite element simulations supply broad parametric coverage, while EIS, microscopy, and inspection measurements correct model discrepancy and support online updating.

Corrosion-specific evidence remains uneven but is expanding. Published studies now address machine learning corrosion prediction [8], stress-corrosion-cracking risk [13], phase-field-assisted prediction for coated magnesium alloys [58], atmospheric and pitting corrosion with PINNs [38,39], subsea pipeline corrosion [40], corrosion fatigue prognosis [87,93], and field or sensor-assisted monitoring [15,94–96]. For organic coatings, neural EIS interpretation, two-stage degradation models, and formulation screening now connect measured barrier response with classification or design decisions [5,97,98]. Table 4 organizes these applications without implying that their metrics are directly comparable.



Calibrate distributions at the field level and at the decision threshold

Figure 5. Hierarchical uncertainty propagation from aleatory inputs and epistemic parameters to degradation fields, decision uncertainty, and value-of-information-based sensing.

Table 4. Representative evidence families and their main methodological implications.

Evidence Family	Representative Sources	Contribution	Remaining Limitation
Corrosion prediction reviews	[8–11]	Establish the breadth of data sources and machine learning tasks	Inconsistent datasets, splits, metrics, and physical closure.
Corrosion-specific PINNs	[37–40]	Move from static regression toward physically constrained temporal prediction	Limited cross-material and cross-environment replication.
Coatings and phase-field hybrids	[58,59,97–99]	Connect interface evolution, service life prediction, and formulation design	Dependence on simulator fidelity and scarce long-duration experiments.
Corrosion fatigue and capacity	[14,87,93]	Link degradation fields to structural response and life	Propagated uncertainty and independent OOD tests remain uncommon.
Inspection and monitoring	[4,6,7,94,100–103]	Demonstrate complementary sensing and data-fusion opportunities	Asset-level independence and sensor failure scenarios require stronger testing.
Materials and inhibitor design	[42,104,105]	Extend machine learning to composition and molecular design	Predicted candidates require mechanistic and experimental confirmation.

3.6. Study-Level Comparison of Corrosion and Coating SciML Studies

The preceding subsections organized the literature by mechanism; this subsection compares the corrosion- and coating-specific implementations study by study. Table 5 records, for each investigation, the material system and exposure environment, the degradation mechanism and input data, the SciML architecture and the physical constraints actually implemented, the validation strategy, and the principal finding together with its main limitation. Each entry carries an evidence level: E when the model was validated against independent experimental or field measurements, S when validation relied on numerical simulations, and S/E for simulation-based studies with partial experimental anchoring. Under this classification, no published corrosion or coating study yet reaches the compiled-physics level of Table 2; that level remains conceptual (C) and defines the research space addressed in Section 4.

Table 5. Study-level comparison of corrosion- and coating-specific SciML investigations. Evidence level: E = validated against independent experimental or field data; S = validated against numerical simulations; S/E = simulation-based with partial experimental anchoring.

Study	System, Material, and Environment	Mechanism and Input Data	Architecture and Physical Constraints	Validation (Level)	Principal Finding and Limitation
[41]	Reinforced concrete; chloride-laden marine and de-icing service	Chloride ingress; measured chloride profiles	PINN; Fickian diffusion residual, inverse diffusivity identification	Published experimental profiles (E)	Recovers apparent diffusivity without meshing; one-dimensional Fickian closure, no binding or convection.
[37]	Steel pipe piles; marine immersion and tidal zones	Corrosion lifetime; environmental and inspection records	PINN with embedded corrosion-kinetics constraints	Field and laboratory data (E)	Physically constrained lifetime prediction; single-site scope, cross-site transfer untested.
[38]	Pipeline steel; buried service	Maximum pit depth; historical inspection data	PINN embedding pit growth kinetics	Field pit depth records (E)	Improves extrapolation of maximum pit depth over data-only baselines; scalar output, no spatial field.
[39]	Carbon and low-alloy steels; multi-site atmospheric exposure	Long-term mass loss; meteorological and pollution histories	Multi-level PINN with stage-wise kinetic laws	Multi-year exposure datasets (E)	Long-horizon prediction beyond the training window; cross-site out-of-distribution replication incomplete.
[40]	Subsea oil and gas pipeline; seawater with operational variables	Internal corrosion rate; operational sensor data	Hybrid PINN–STAN (attention plus physics residuals)	Operational pipeline data (E)	Accuracy gain over purely data-driven baselines; attention terms weakly interpretable, single asset.
[93]	Aluminum alloys; corrosion fatigue	Crack growth rate; multifactorial fatigue tests	PINN embedding Paris-type crack growth law	Laboratory fatigue datasets (E)	Multifactor crack growth prediction; deterministic, uncertainty not propagated.
[87]	High-strength steel with non-uniform corrosion	Fatigue life; S–N data and corrosion morphology	Adaptive Bayesian PINN	Laboratory fatigue and morphology data (E)	Calibrated intervals on fatigue life; small specimen set, coverage testing limited.
[42]	Ferritic/martensitic steels; lead–bismuth eutectic	Oxide growth kinetics; composition and exposure data	Physics-guided network with monotonic oxide growth priors	Published corrosion rate experiments (E)	Composition optimization consistent with kinetics; extrapolation beyond the alloy family unverified.
[58]	Coated biodegradable Mg alloys; simulated body fluid	Interface recession; phase-field simulation ensembles	Phase-field simulations coupled to a machine-learned surrogate	Simulations with partial experimental anchoring (S/E)	Order-of-magnitude acceleration for coated-Mg screening; validation partly circular with respect to the simulator.
[99]	Generic metal–electrolyte pitting	Phase-field corrosion equations	Staggered, hard-constrained PINN	Numerical phase-field reference solutions (S)	Staggering mitigates stiff coupled optimization; preprint status, no experimental validation.
[59]	Zn–Al–Mg-coated high-strength steel wires	Stress-accelerated dissolution; mechanically coupled phase field	Mechanistic phase-field model (machine learning-ready)	Simulations benchmarked to published tests (S/E)	Quantifies stress acceleration in coated wires; parameter identification remains manual.
[97]	Organic coatings on steel; multi-environment field exposure	Two-stage environment–property–failure mapping	Physics-guided two-stage machine learning (no PDE residual)	Field exposure datasets (E)	Interpretable degradation staging; correlational closure without governing equation constraint.

Table 5. *Cont.*

Study	System, Material, and Environment	Mechanism and Input Data	Architecture and Physical Constraints	Validation (Level)	Principal Finding and Limitation
[5]	Organic coatings; laboratory immersion	EIS spectra interpretation	Neural classification and regression on impedance data (physics-guided)	Laboratory EIS measurements (E)	Automates impedance interpretation; depends on the measurement model, no temporal dynamics.
[98]	Anticorrosive coating formulations; accelerated testing	Formulation–performance mapping; historic test data	Machine learning screening of historic industrial data	Historic experimental records (E)	Suggests novel formulation candidates; mechanistic and experimental confirmation pending.

Table 6 complements this comparison by summarizing, for each model family, the materials and corrosive environments in which it has actually been demonstrated, because corrosion is too heterogeneous a phenomenon for any single architecture to claim general applicability.

Table 6. Demonstrated applicability of SciML model families across materials and corrosive environments, with the principal caution governing transfer to new conditions.

Model Family	Materials Demonstrated	Environments Demonstrated	Transfer Caution
PINN (forward and inverse) [37–41]	Carbon and pipeline steels; reinforced concrete; steel piles	Marine immersion and tidal; multi-site atmospheric; buried soil; subsea operational	Retraining is required per geometry and parameter set; cross-environment transfer is largely untested.
Phase-field PINN and phase-field/ML hybrids [58,59,99]	Coated Mg alloys; Zn–Al–Mg coated steel wires; generic pitting metals	Simulated body fluid; stress-assisted laboratory conditions; synthetic electrolytes	Validation is dominated by simulators; independent experimental interface tracking is still missing.
Physics-guided regression and kinetics-constrained networks [5,42,97,98]	Ferritic/martensitic steels; organic coatings on steel; coating formulations	High-temperature liquid metal; field atmospheric exposure; accelerated laboratory testing	Physical knowledge enters only through variables or monotonic priors; no governing equation closure is enforced.
Bayesian and ensemble SciML [86–89]	High-strength steel under corrosion fatigue	Laboratory corrosion fatigue testing	Interval calibration has been verified only on small specimen sets.
Neural operators [74–77]	No corrosion-specific deployment published; generic parametric PDE benchmarks	—	Grid dependence and geometric transfer to evolving pits remain unproven in corrosion settings.
GNN and molecular/crystal graph models [105–107]	Alloy and inhibitor chemistry via property prediction	Composition screening rather than exposure	Conservation and irreversibility are not automatic; coupling to field-scale degradation is unrealized.

Three comparative observations follow. First, the studies converge on one finding: embedding kinetic or transport structure improves extrapolation beyond the training window relative to data-only baselines, a result replicated independently for chloride ingress, atmospheric mass loss, pit depth, and liquid-metal oxide growth [38,39,41,42]. Second, the studies disagree, or cannot be compared, on the magnitude of that benefit: reported metrics mix coefficients of determination, root-mean-square errors on different scales, and heterogeneous data splits; equal-budget baselines are rare; and the phase-field surrogates report acceleration factors that exclude training cost [58,99]. Whether attention-based or hybrid variants outperform a well-tuned standard PINN under an equal computational budget is presently an open question [40]. Third, the unresolved limitations are structural rather than incremental: no study simultaneously demonstrates discrete conservation, nonnegative dissipation, and calibrated uncertainty under out-of-distribution evaluation; no cross-material replication of a single trained model has been published; and the organic coating studies remain at the physics-guided level, so governing equation closure for polymeric barriers is still an open problem. These three deficits motivate the methodological requirements of Section 4.

3.7. Relevance to Coating Systems and Industrial Practice

For the coatings community, the value of SciML depends on how directly a model addresses specific coating classes, exposure environments, and measurable degradation

mechanisms. Table 7 summarizes this mapping. Barrier organic systems (epoxy intermediates and polyurethane topcoats) degrade through water uptake, blistering, cathodic delamination, and under-film corrosion, processes quantified primarily by EIS, capacitance tracking, adhesion testing, and scanning probes such as SVET and SKP [1,2,108]. Zinc-rich and sacrificial systems fail by galvanic depletion, for which phase-field galvanic models exist but SciML applications remain absent [49]. Metallic Zn–Al–Mg coatings, conversion and anodized layers, biodegradable magnesium coatings, and inhibitor-doped smart coatings each carry distinct mechanisms and measurement practices, and the table shows that SciML coverage across these classes is highly uneven [12,58,59,104,105].

Table 7. Coating classes, dominant degradation mechanisms, characteristic measurements, and current SciML status.

Coating Class	Dominant Degradation Mechanisms	Characteristic Measurements	Current SciML Status
Barrier organic coatings (epoxy, polyurethane)	Water uptake, blistering, cathodic delamination, under-film corrosion	EIS, coating capacitance, SKP, adhesion, salt spray	Physics-guided degradation staging and neural EIS interpretation [4,5,97].
Zinc-rich primers and sacrificial systems	Galvanic depletion, zinc consumption, loss of electrical contact	Open-circuit potential, SVET, mass loss	Essentially unexplored by SciML; mechanistic phase-field galvanic models available [49].
Metallic Zn–Al–Mg coatings	Stress-assisted dissolution, cut-edge attack	Accelerated wire tests, cross-sectional microscopy	Mechanistic phase field with machine learning coupling incipient [59].
Conversion, anodized, and PEO layers	Loss of pore sealing, substrate pitting	EIS, potentiodynamic polarization, electron microscopy	Data-driven property prediction only; governing equation closure missing [8,104].
Biodegradable Mg coatings	Controlled interface recession, hydrogen evolution	Hydrogen collection, mass loss, in vitro EIS	Phase-field simulation with learned surrogates [58].
Inhibitor-doped smart coatings	Inhibitor leaching, triggered release, self-healing	EIS, local electrochemical probes, release assays	Molecular screening of inhibitors; mechanistic release models not yet coupled [12,105].

Exposure environments should be reported in the shared vocabulary of corrosivity classification: the atmospheric categories of ISO 9223 and the C1–CX and Im1–Im4 classes of ISO 12944-2 [109,110]. Mapped onto that vocabulary, the studies of Table 5 cover multi-site atmospheric service [39], marine tidal and immersion zones [37], subsea operation [40], buried pipelines [38], physiological media [58], and heavy-liquid-metal service outside the ISO classes altogether [42]. We recommend that future corrosion SciML studies state the corrosivity class of both training and test exposures explicitly, because environment-blind reporting is one reason the current evidence base resists comparison.

Regarding data resources, long-duration coating datasets with documented exposure histories remain the scarcest asset. Image-based inspection corpora exist and support convolutional and unsupervised staging models [35,95,100,111], and complementary sensing streams—acoustic emission, embedded RFID, cathodic protection records, and field computer vision—have been demonstrated individually [94,96,101,102]. However, the community-level call to share corrosion data in curated, reusable form remains only partially answered a decade after it was made [112]. The practical barriers to industrial implementation are equally concrete: proprietary formulations restrict disclosure of coating chemistry; accelerated tests correlate imperfectly with field performance, so models trained on salt spray data inherit that uncertainty; qualification standards do not yet recognize learned models as evidence; and certified deployment requires the versioning, limits-of-use documentation, and fallback behavior discussed in Section 4. These constraints, more than architectural novelty, currently determine whether a coating SciML model can leave the laboratory.

3.8. Topology and Validation Beyond Random Splits

Pixelwise error and average thickness loss do not distinguish two damage fields with equal lost volume but different connectivity. Betti numbers, Euler characteristic, and persistence diagrams provide complementary measures of pit nucleation, isolated cavities, connected attack, and percolating paths [113,114]. Differentiable losses based on persistent homology show how such descriptors can enter model training rather than remain post hoc diagnostics [115]. Figure 6 illustrates these descriptors using synthetic morphology. The figure is explanatory and should not be read as a measured coating history.




Figure 6. Conceptual topology-sensitive diagnostics for localized degradation: (a) persistence diagram for connected components and loops; (b) evolution of Betti numbers and Euler characteristic during nucleation and coalescence. Values are synthetic and illustrate metrics rather than experimental results.

A credible validation protocol must challenge the physical assumptions that the model claims to encode. Table 8 lists minimum tests. Random record-level splitting may remain useful for debugging, but it cannot be the sole evidence for deployment when neighboring measurements share a specimen, exposure history, simulation, or asset.

Table 8. Minimum validation protocol proposed for corrosion and coating SciML.

Challenge	Purpose	Minimum Reporting
Leave one environment out	Test transfer across humidity, salinity, temperature, or loading histories	Per-environment error, worst-case error, and uncertainty coverage.
Leave one material or coating out	Test chemical and microstructural transfer	Material identifiers, composition ranges, calibration data, and failure cases.
Temporal extrapolation	Test long-horizon stability and memory	Horizon-dependent error, drift, mass/charge balance, and free energy rate.
Geometry or mesh transfer	Test pit, defect, and discretization generalization	Field error, maximum depth, lost volume, and topology-sensitive error.
Sensor degradation	Test missing, biased, or noisy measurements	Calibration, coverage, robustness curves, and recovery strategy.
Physics ablation	Identify the value of each encoded constraint	Equal-budget baselines and removal of conservation, memory, topology, and UQ terms one at a time.
Solver and experiment comparison	Separate emulation from physical validation	Analytical, FEM/phase-field, and independent experimental comparisons with documented discrepancy.
Computational budget	Determine whether acceleration is operationally meaningful	Training cost, inference latency, memory, hardware, and accuracy–cost trade-off.

4. A Thermo-Fractional Graph Phase-Field Neural Operator Research Framework

4.1. Motivation and Scope

The evidence reviewed above, and in particular the structural deficits recorded in Tables 5 and 6, does not identify a single validated architecture that simultaneously resolves non-Fickian transport, electrochemical reactions, interface evolution, irregular geometry, thermodynamic irreversibility, and calibrated uncertainty. The proposed Thermo-Fractional Graph Phase-Field Neural Operator (TFG–PFNO) is intended to make this integration explicit. It is a research framework and falsifiable design hypothesis, not a model whose superiority is established by the conceptual comparisons in this paper.

The framework addresses five recurring gaps. First, a fractional memory mechanism represents transport histories that cannot be summarized by an instantaneous effective diffusivity. Second, a graph encoder represents coating microstructure, defects, interfaces, meshes, and sensor networks without assuming a fixed Cartesian topology. Third, a phase-field head represents the evolution and coalescence of degraded regions. Fourth, a metriplectic or GENERIC-compatible latent dynamics separates reversible transport from irreversible dissipation, following structure-preserving, thermodynamics-informed, and GENERIC-informed learning developments [116–119]. The underlying metriplectic and GENERIC formalisms were established for dissipative continuum systems well before their neural adaptations [120,121], so the separation proposed here inherits a mature nonequilibrium thermodynamic foundation rather than an architectural convenience. Fifth, probabilistic inference propagates uncertainty from inputs and parameters to failure-relevant outputs.

Figure 7 shows the proposed information flow. Molecular or microstructural descriptors and interface or sensor graphs are encoded into an equivariant latent representation. A fractional memory block processes environmental and loading histories. Reversible and dissipative latent operators then evolve the state before phase-field and probabilistic decoders produce concentrations, potential, damage, interface state, and uncertainty.




Figure 7. Proposed TFG–PFNO research framework. Graph encoders represent material chemistry and irregular geometry; a fractional block represents memory; reversible and dissipative operators evolve the latent state; and phase-field and probabilistic heads predict degradation fields and uncertainty. The architecture is conceptual and requires implementation, ablation, and experimental validation.

4.2. Structured Latent Dynamics

Let a denote static material and geometric attributes, $\mathbf{s}_{0:t}$ the environmental and loading history, and $\mathcal{G}_\Omega$ the graph representation of the coating–electrolyte–substrate system. The encoder produces an initial latent state,

$$\mathbf{z}_0 = \mathcal{E}_{G,\theta}(a, \mathcal{G}_\Omega, \mathbf{s}_0). \quad (13)$$

Its evolution is represented schematically as

$${}^C D_t^{\alpha_\theta} \mathbf{z} = \mathbf{L}_\theta(\mathbf{z}) \nabla E_\theta(\mathbf{z}) + \mathbf{M}_\theta(\mathbf{z}) \nabla S_\theta(\mathbf{z}) + \mathcal{B}_\theta \mathbf{s}(t), \quad (14)$$

with

$$\mathbf{L}_\theta^\top = -\mathbf{L}_\theta, \quad \mathbf{M}_\theta = \mathbf{B}_\theta \mathbf{B}_\theta^\top + \epsilon \mathbf{I} \succeq 0. \quad (15)$$

The antisymmetric operator $\mathbf{L}_\theta$ carries reversible transport, whereas the positive semidefinite operator $\mathbf{M}_\theta$ generates entropy-producing or free-energy-reducing dynamics, depending on the adopted thermodynamic potential and sign convention. Degeneracy conditions must prevent the reversible channel from producing entropy and the dissipative channel from altering invariants that should remain conserved. These conditions cannot be assumed from Equation (15) alone; they must be stated and tested for the selected state variables.

The decoder maps the latent trajectory and query coordinates to physical predictions,

$$(\hat{\mathbf{c}}, \hat{\psi}, \hat{\phi}, \hat{d}, \hat{\Sigma}) = \mathcal{D}_\theta(\mathbf{z}, \mathbf{x}, t), \quad (16)$$

where $\hat{\Sigma}$ represents predictive uncertainty or the parameters of a predictive distribution. The phase-field variables may be decoded on a mesh or graph and then used to update connectivity. A conservative implementation should specify how node or cell volumes, interface areas, and source terms change when the degraded region evolves.

The fractional order α_θ should be constrained to a physically meaningful interval and subjected to identifiability tests. A variable-order model may be justified when damage or saturation changes the transport regime, but it increases confounding between memory, diffusivity, and reaction kinetics. Recovery experiments with synthetic truth, followed by independent coating exposures, are therefore necessary before interpreting the learned fractional order mechanistically.

4.3. Training Strategy

A practical implementation should proceed in stages rather than optimize the complete model from random initialization. First, the graph encoder and operator backbone can be pretrained on physically diverse simulation ensembles. Second, hard constraints and positive parameterizations can be activated for quantities with unambiguous signs or boundary values. Third, electrochemical, phase-field, and mechanical blocks can be coupled progressively using alternating or curriculum schedules. Finally, the probabilistic head can be calibrated on held-out experimental or field data.

The training objective should combine dimensionless residuals with explicit probabilistic and geometric terms,

$$\begin{aligned} \mathcal{L}_{\text{TFG-PFNO}} = & \lambda_{\text{obs}} \mathcal{L}_{\text{obs}} + \lambda_{\text{mass}} \tilde{\mathcal{L}}_{\text{mass}} + \lambda_{\text{charge}} \tilde{\mathcal{L}}_{\text{charge}} + \lambda_{\text{phase}} \tilde{\mathcal{L}}_{\text{phase}} \\ & + \lambda_{\text{mech}} \tilde{\mathcal{L}}_{\text{mech}} + \lambda_{\text{thermo}} \tilde{\mathcal{L}}_{\text{thermo}} + \lambda_{\text{topo}} \mathcal{L}_{\text{topo}} + \lambda_{\text{cal}} \mathcal{L}_{\text{cal}}. \end{aligned} \quad (17)$$

The observation term must account for the relevant measurement operator rather than compare latent fields directly with heterogeneous sensor outputs. The topological term can compare persistence diagrams or differentiable approximations of connectivity. The calibration term can use a proper scoring rule, while final evaluation should report reliability and interval coverage in addition to mean error.

Adaptive weighting may stabilize optimization, but the hierarchy of constraints should be explicit. Mass conservation, charge balance, positivity, and the adopted thermodynamic inequality should not be allowed to disappear because their gradients are temporarily small or antagonistic. Conversely, over-weighting an uncertain closure rela-

tion can force systematic bias. The loss therefore requires both numerical balancing and epistemic ranking: mature physical laws receive stronger enforcement than speculative constitutive assumptions.

4.4. Benchmark and Falsification Plan

The proposed framework should be compared with simpler alternatives under equal computational budgets. At minimum, the benchmark set should include a data-only neural network, a standard PINN, a hard-constrained or variational PINN, a neural operator, a geometry-aware GNN, and the full TFG–PFNO. Ablations should remove the fractional, graph, phase-field, thermodynamic, and probabilistic components individually. A complex model is justified only if its improvement appears on the specific challenge that the removed component was designed to address.

The first benchmark tier should use analytical or high-resolution numerical problems with known parameters and controlled noise. These tests can establish convergence, parameter recovery, conservation, and computational scaling. The second tier should use accelerated coating exposures with independent measurements of water uptake, EIS, surface or cross-sectional morphology, and mass loss. The third tier should use long-duration or field data with changing environments and missing sensors. At each tier, training and test cases should be separated by physical unit, not by neighboring samples.

Failure criteria should be defined before model comparison. Examples include negative concentrations, increasing damage reversal, violation of mass or charge balance beyond tolerance, nonphysical free energy production, collapse of interval coverage, loss of connectivity fidelity, and excessive inference latency. Reporting only the best seed or the lowest field error would not establish reliability. The benchmark should publish all seeds, hyperparameter ranges, unsuccessful runs, and code required to reproduce the principal figures.

4.5. Implications for Corrosion Digital Twins

A corrosion digital twin is operationally useful only when it connects state estimation to a decision. The posterior distribution predicted by TFG–PFNO should therefore be propagated to remaining capacity, failure probability, inspection interval, or expected maintenance loss. Sensor placement can then be framed through value of information: a measurement is valuable when it reduces decision-relevant uncertainty enough to justify its cost.

Industrial deployment also requires explicit limits of use. The twin should record the material chemistries, coating systems, environments, geometries, and loading regimes represented in training. When an input lies outside those limits or predictive uncertainty exceeds a validated threshold, the system should enter a conservative fallback mode, request inspection, or revert to a trusted numerical or rule-based model. Versioned data, equations, calibration records, and model cards are therefore part of the scientific method rather than administrative extras.

Finally, atomistic or quantum information should enter through documented constitutive closure, not through an undefined claim of “quantum physics in the loss”. Graph models for crystals and molecules and equivariant interatomic potentials can estimate material properties, adsorption energies, or activation barriers that inform distributions of reaction and transport parameters [106,107,122–124]. The coarse-graining procedure must preserve stoichiometry, relevant affinities, and uncertainty. This modular route is more transparent than training one end-to-end network across electronic and structural scales without an identifiable transfer operator.

5. Conclusions

SciML is moving corrosion and coating degradation prediction from isolated correlations toward learned fields, inverse parameters, and solution operators. The transition is scientifically important, but the evidence does not support treating every PDE-regularized network as a reliable physical model. Current studies remain fragmented across electrochemistry, coating transport, phase-field evolution, structural mechanics, geometry, and uncertainty. The principal need is therefore not a deeper network in isolation, but a mathematical and validation structure that makes conservation, irreversibility, and transfer explicit.

Five conclusions follow from the reviewed literature. First, composite losses must be nondimensionalized and evaluated by mechanism; total loss is not a physical metric. Second, mature constraints such as positivity, internal flux conservation, and positive-semidefinite dissipation should be encoded structurally when feasible. Third, anomalous transport requires identifiable memory models and long-horizon validation rather than an unconstrained effective coefficient. Fourth, pit and defect connectivity calls for geometry-aware representations and topology-sensitive metrics in addition to pixelwise or average-field errors. Fifth, uncertainty must be calibrated and propagated to structural risk, inspection planning, and maintenance decisions.

The study-level comparison of Table 5 sharpens these conclusions into evidence-based gaps. Experimentally validated demonstrations are concentrated on scalar or low-dimensional targets—pit depth, corrosion rate, lifetime, and fatigue life [37–39,93]—whereas full-field interface prediction has been validated almost exclusively against the simulators that generated the training data [58,99]. Organic-coating applications remain at the physics-guided level, without governing equation closure for polymeric barriers [5,97]. Calibrated uncertainty has been demonstrated in corrosion fatigue only on small specimen sets [87], and no published study yet combines discrete conservation, nonnegative dissipation, and calibrated uncertainty under out-of-distribution evaluation. These are the gaps that any credible next generation of corrosion SciML must close.

The proposed TFG–PFNO framework combines a graph encoder, fractional memory, phase-field interface evolution, structured reversible and dissipative dynamics, and probabilistic prediction. Its value is presently conceptual: each module corresponds to a documented limitation, and the complete architecture still requires implementation, equal-budget ablation, and independent experimental validation; it should therefore be read as a falsifiable synthesis of the gaps above rather than as a recommendation for immediate deployment.

Future work in this field should advance along six concrete directions. (1) Open multiscale benchmarks that pair accelerated coating exposures with documented environmental histories, specimen- and asset-independent partitions, and corrosivity classes reported per ISO 9223 and ISO 12944-2 [109,110,112]. (2) Experimentally validated structure preservation: implementations whose conservation and dissipation properties are demonstrated on measured pit growth and delamination, not only on synthetic phase-field solutions [57,99]. (3) Identifiable memory: fractional-order transport models for coating water uptake whose parameters are recovered from independent gravimetric and EIS data before mechanistic interpretation [47,108]. (4) Governing equation closure for organic coatings, connecting EIS observables to transport and delamination fields through explicit measurement operators [4,5]. (5) Decision-grade uncertainty: calibration-tested posteriors propagated to inspection intervals and remaining-capacity estimates, with simulator discrepancy modeled explicitly [87,92]. (6) Accelerated-to-field transfer, treating laboratory and field exposures as different fidelities within a discrepancy-aware multifidelity framework [90,91].

For the coatings community, the immediate priorities are open multiscale benchmarks; specimen- and asset-independent data partitions; standardized conservation, topology, and calibration metrics; and transparent reporting of failed as well as successful training runs. These practices will determine whether corrosion digital twins become traceable engineering tools or remain high-accuracy interpolators with uncertain behavior outside the laboratory conditions used to train them.

Author Contributions: Conceptualization, L.R.-V. and J.G.; methodology, L.R.-V.; formal analysis, L.R.-V. and P.M.; investigation, L.R.-V., P.M. and A.P.; writing—original draft preparation, L.R.-V.; writing—review and editing, P.M., A.P. and J.G.; visualization, L.R.-V.; supervision, A.P. and J.G. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Acknowledgments: During manuscript preparation, the authors used OpenAI Codex (GPT-5) for English translation, editorial restructuring, reference format consistency checks, and figure format normalization. The authors reviewed the resulting text and graphics and remain fully responsible for the content.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

CNN	Convolutional neural network
DOI	Digital object identifier
EIS	Electrochemical impedance spectroscopy
FEM	Finite element method
FNO	Fourier neural operator
GENERIC	General Equation for Nonequilibrium Reversible-Irreversible Coupling
GNN	Graph neural network
ILI	In-line inspection
ISO	International Organization for Standardization
ML	Machine learning
NACE	National Association of Corrosion Engineers
OOD	Out of distribution
PDE	Partial differential equation
PEO	Plasma electrolytic oxidation
PINN	Physics-informed neural network
RFID	Radio-frequency identification
SciML	Scientific machine learning
SDF	Signed distance function
SHM	Structural health monitoring
SKP	Scanning Kelvin probe
STAN	Spatio-temporal attention network
SVET	Scanning vibrating electrode technique
TFG-PFNO	Thermo-fractional graph phase-field neural operator
UQ	Uncertainty quantification

References

1. Sørensen, P.A.; Kiil, S.; Dam-Johansen, K.; Weinell, C.E. Anticorrosive coatings: A review. *J. Coat. Technol. Res.* **2009**, *6*, 135–176. [\[CrossRef\]](#)
2. Grundmeier, G.; Schmidt, W.; Stratmann, M. Corrosion protection by organic coatings: Electrochemical mechanism and novel methods of investigation. *Electrochim. Acta* **2000**, *45*, 2515–2533. [\[CrossRef\]](#)
3. Lyon, S.B.; Bingham, R.; Mills, D.J. Advances in corrosion protection by organic coatings: What we know and what we would like to know. *Prog. Org. Coat.* **2017**, *102*, 2–7. [\[CrossRef\]](#)
4. Policastro, S.A.; Anderson, R.M.; Hangarter, C.M.; Arcari, A.; Iezzi, E.B. Incorporating Physics-Based Models into Equivalent Circuit Analysis of EIS Data from Organic Coatings. *Coatings* **2023**, *13*, 1285. [\[CrossRef\]](#)
5. Bongiorno, V.; Michailidou, E.; Curioni, M. Towards the automated interpretation of impedance spectra from organic coatings using neural networks. *Corros. Sci.* **2024**, *230*, 111922. [\[CrossRef\]](#)
6. Bongiorno, V.; Michailidou, E.; Curioni, M.; Lyon, S. Unsupervised classification of organic coating performance using electrochemical impedance data. *J. Appl. Electrochem.* **2025**, *55*, 1947–1962. [\[CrossRef\]](#)
7. Martinez, S.; Hudec, B.; Martinez, I.; Šoić, I. Unsupervised clustering of EIS spectra for distinguishing degradation states of polymer coatings. *Prog. Org. Coat.* **2026**, *215*, 110082. [\[CrossRef\]](#)
8. Coelho, L.B.; Zhang, D.; Van Ingelgem, Y.; Steckelmacher, D.; Nowe, A.; Terryn, H. Reviewing machine learning of corrosion prediction in a data-oriented perspective. *npj Mater. Degrad.* **2022**, *6*, 8. [\[CrossRef\]](#)
9. Imran, M.M.H.; Jamaludin, S.; Ayob, A.F.M.; Ali, A.A.I.M.; Ahmad, S.Z.A.S.; Akhbar, M.F.A.; Suhrab, M.I.R.; Zainal, N.; Norzeli, S.M.; Mohamed, S.B. Application of artificial intelligence in marine corrosion prediction and detection. *J. Mar. Sci. Eng.* **2023**, *11*, 256. [\[CrossRef\]](#)
10. Hussain, M.; Zhang, T.; Chaudhry, M.; Jamil, I.; Kausar, S.; Hussain, I. Review of prediction of stress corrosion cracking in gas pipelines using machine learning. *Machines* **2024**, *12*, 42. [\[CrossRef\]](#)
11. Guo, R.; Rong, J. Machine learning in corrosion electrochemistry: A critical review of prediction and corrosion-type identification. *Int. J. Electrochem. Sci.* **2026**, *21*, 101324. [\[CrossRef\]](#)
12. Mubeen, M.; Khalid, S.; Tabish, M.; Bibi, F.; Ilyas, H.M.A.; Chen, S.; Yasin, G.; Zhao, J.; Yang, C. Performance evolution, predictive modeling and smart nanocomposite strategies for organic corrosion protection coatings. *Corros. Commun.* **2026**, Advance online publication. [\[CrossRef\]](#)
13. Alamri, A.H. Application of machine learning to stress corrosion cracking risk assessment. *Egypt. J. Pet.* **2022**, *31*, 11–21. [\[CrossRef\]](#)
14. Wang, Q.; Lu, H. Machine learning methods for predicting residual strength in corroded oil and gas steel pipes. *npj Mater. Degrad.* **2025**, *9*, 30. [\[CrossRef\]](#)
15. Olawole, O.O.; Gbadeyan, O.J.; Deigh, C.; Sijuade, M.F. Advanced sensor systems and machine learning for pipeline integrity management: A review of corrosion monitoring and prediction strategies. *npj Mater. Degrad.* **2026**, *10*, 48. [\[CrossRef\]](#)
16. Lagaris, I.E.; Likas, A.; Fotiadis, D.I. Artificial neural networks for solving ordinary and partial differential equations. *IEEE Trans. Neural Netw.* **1998**, *9*, 987–1000. [\[CrossRef\]](#) [\[PubMed\]](#)
17. Sirignano, J.; Spiliopoulos, K. DGM: A deep learning algorithm for solving partial differential equations. *J. Comput. Phys.* **2018**, *375*, 1339–1364. [\[CrossRef\]](#)
18. Karniadakis, G.E.; Kevrekidis, I.G.; Lu, L.; Perdikaris, P.; Wang, S.; Yang, L. Physics-informed machine learning. *Nat. Rev. Phys.* **2021**, *3*, 422–440. [\[CrossRef\]](#)
19. Cuomo, S.; Di Cola, V.S.; Giampaolo, F.; Rozza, G.; Raissi, M.; Piccialli, F. Scientific machine learning through physics-informed neural networks: Where we are and what is next. *J. Sci. Comput.* **2022**, *92*, 88. [\[CrossRef\]](#)
20. Cai, S.; Mao, Z.; Wang, Z.; Yin, M.; Karniadakis, G.E. Physics-informed neural networks (PINNs) for fluid mechanics: A review. *Acta Mech. Sin.* **2021**, *37*, 1727–1738. [\[CrossRef\]](#)
21. Lawal, Z.K.; Yassin, H.; Lai, D.T.C.; Che Idris, A. Physics-Informed Neural Network (PINN) Evolution and Beyond: A Systematic Literature Review and Bibliometric Analysis. *Big Data Cogn. Comput.* **2022**, *6*, 140. [\[CrossRef\]](#)
22. Raissi, M.; Perdikaris, P.; Karniadakis, G.E. 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. [\[CrossRef\]](#)
23. Lu, L.; Meng, X.; Mao, Z.; Karniadakis, G.E. DeepXDE: A Deep Learning Library for Solving Differential Equations. *SIAM Rev.* **2021**, *63*, 208–228. [\[CrossRef\]](#)
24. Krishnapriyan, A.; Gholami, A.; Zhe, S.; Kirby, R.; Mahoney, M.W. Characterizing possible failure modes in physics-informed neural networks. *Adv. Neural Inf. Process. Syst.* **2021**, *34*, 26548–26560. [\[CrossRef\]](#)
25. Wang, S.; Teng, Y.; Perdikaris, P. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM J. Sci. Comput.* **2021**, *43*, A3055–A3081. [\[CrossRef\]](#)
26. Wang, S.; Yu, X.; Perdikaris, P. When and why PINNs fail to train: A neural tangent kernel perspective. *J. Comput. Phys.* **2022**, *449*, 110768. [\[CrossRef\]](#)

27. Wang, S.; Sankaran, S.; Perdikaris, P. Respecting causality for training physics-informed neural networks. *Comput. Methods Appl. Mech. Eng.* **2024**, *421*, 116813. [[CrossRef](#)]
28. Yu, J.; Lu, L.; Meng, X.; Karniadakis, G.E. Gradient-enhanced physics-informed neural networks for forward and inverse PDE problems. *Comput. Methods Appl. Mech. Eng.* **2022**, *393*, 114823. [[CrossRef](#)]
29. Meng, X.; Li, Z.; Zhang, D.; Karniadakis, G.E. PPINN: Parareal physics-informed neural network for time-dependent PDEs. *Comput. Methods Appl. Mech. Eng.* **2020**, *370*, 113250. [[CrossRef](#)]
30. Xiang, Z.; Peng, W.; Liu, X.; Yao, W. Self-adaptive loss balanced physics-informed neural networks. *Neurocomputing* **2022**, *496*, 11–34. [[CrossRef](#)]
31. McClenny, L.; Braga-Neto, U. Self-adaptive physics-informed neural networks. *J. Comput. Phys.* **2023**, *474*, 111722. [[CrossRef](#)]
32. Tang, K.; Wan, X.; Yang, C. DAS-PINNs: A deep adaptive sampling method for solving high-dimensional partial differential equations. *J. Comput. Phys.* **2023**, *476*, 111868. [[CrossRef](#)]
33. Mao, Z.; Meng, X. Physics-informed neural networks with residual/gradient-based adaptive sampling methods for solving partial differential equations with sharp solutions. *Appl. Math. Mech.* **2023**, *44*, 1069–1084. [[CrossRef](#)]
34. Koch, G.; Varney, J.; Thompson, N.; Moghissi, O.; Gould, M.; Payer, J. *International Measures of Prevention, Application, and Economics of Corrosion Technologies Study (IMPACT)*; Technical Report; NACE International: Houston, TX, USA, 2016.
35. Nash, W.; Drummond, T.; Birbilis, N. A review of deep learning in the study of materials degradation. *npj Mater. Degrad.* **2018**, *2*, 37. [[CrossRef](#)]
36. Raissi, M.; Yazdani, A.; Karniadakis, G.E. Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations. *Science* **2020**, *367*, 1026–1030. [[CrossRef](#)] [[PubMed](#)]
37. Feng, Z.; Jiang, Z.; Yan, C. Prediction of corrosion lifetime of marine environment steel pipe piles based on physics-informed neural network. *Discov. Appl. Sci.* **2025**, *8*, 136. [[CrossRef](#)]
38. Hu, Q.; Zhang, Z.; Wang, F.; Su, Z.; Zhou, J. A physics-informed neural network method for predicting maximum pitting corrosion depth in pipelines. *Process Saf. Environ. Prot.* **2026**, *208*, 108520. [[CrossRef](#)]
39. Chen, Q.; Wang, H.; Ma, X. Multi-level physics-informed neural networks (ML-PINN): A framework for long-term atmospheric corrosion prediction. *Corros. Sci.* **2026**, *267*, 113902. [[CrossRef](#)]
40. Zhong, Y.; Dong, X.; Li, Y.; He, S. A novel PINN-STAN model for predicting corrosion rate in subsea oil and gas pipelines. *Ocean Eng.* **2026**, *350*, 124260. [[CrossRef](#)]
41. Wan, Y.; Zheng, W.; Wang, Y. Identification of chloride diffusion coefficient in concrete using physics-informed neural networks. *Constr. Build. Mater.* **2023**, *393*, 132049. [[CrossRef](#)]
42. Feng, S.; Sun, X.; Chen, G.; Chen, X. The lead-bismuth eutectic corrosion rate prediction and composition optimization of ferritic/martensitic steels by physics-guided neural network. *Eng. Appl. Artif. Intell.* **2025**, *141*, 109807. [[CrossRef](#)]
43. Metzler, R.; Klafter, J. The random walk's guide to anomalous diffusion: A fractional dynamics approach. *Phys. Rep.* **2000**, *339*, 1–77. [[CrossRef](#)]
44. Mainardi, F. *Fractional Calculus and Waves in Linear Viscoelasticity*; Imperial College Press: London, UK, 2010. [[CrossRef](#)]
45. Podlubny, I. *Fractional Differential Equations*; Academic Press: San Diego, CA, USA, 1999.
46. Caputo, M. Linear Models of Dissipation whose Q is almost Frequency Independent–II. *Geophys. J. Int.* **1967**, *13*, 529–539. [[CrossRef](#)]
47. Pang, G.; Lu, L.; Karniadakis, G.E. fPINNs: Fractional physics-informed neural networks. *SIAM J. Sci. Comput.* **2019**, *41*, A2603–A2626. [[CrossRef](#)]
48. Mai, W.; Soghrati, S. A phase field model for simulating the pitting corrosion. *Corros. Sci.* **2016**, *110*, 157–166. [[CrossRef](#)]
49. Mai, W.; Soghrati, S. New phase field model for simulating galvanic and pitting corrosion processes. *Electrochim. Acta* **2018**, *260*, 290–304. [[CrossRef](#)]
50. Ansari, T.Q.; Xiao, Z.; Hu, S.; Li, Y.; Luo, J.L.; Shi, S.Q. Phase-field model of pitting corrosion kinetics in metallic materials. *npj Comput. Mater.* **2018**, *4*, 38. [[CrossRef](#)]
51. Chadwick, A.F.; Stewart, J.A.; Enrique, R.A.; Du, S.; Thornton, K. Numerical Modeling of Localized Corrosion Using Phase-Field and Smoothed Boundary Methods. *J. Electrochem. Soc.* **2018**, *165*, C633–C646. [[CrossRef](#)]
52. Sahu, S.; Frankel, G.S. Phase Field Modeling of Crystallographic Corrosion Pits. *J. Electrochem. Soc.* **2022**, *169*, 020557. [[CrossRef](#)]
53. Cui, C.; Ma, R.; Martínez-Pañeda, E. Electro-chemo-mechanical phase field modeling of localized corrosion: Theory and COMSOL implementation. *Eng. Comput.* **2023**, *39*, 3877–3894. [[CrossRef](#)]
54. Zhi, H.; Dong, P.; Li, K.; Gao, L.; Zhou, W.; Zhang, H. Phase field study of pitting corrosion: Electrochemical reactions and temperature dependence. *Comput. Mater. Sci.* **2024**, *244*, 113251. [[CrossRef](#)]
55. Frankel, G.S. Pitting corrosion of metals: A review of the critical factors. *J. Electrochem. Soc.* **1998**, *145*, 2186–2198. [[CrossRef](#)]
56. Macdonald, D.D. The point defect model for the passive state. *J. Electrochem. Soc.* **1992**, *139*, 3434–3449. [[CrossRef](#)]
57. Cui, C.; Ma, R.; Martínez-Pañeda, E. A phase field formulation for dissolution-driven stress corrosion cracking. *J. Mech. Phys. Solids* **2021**, *147*, 104254. [[CrossRef](#)]

58. Ma, S.; Zhang, D.; Zhang, P.; Markert, B. Rapid prediction of the corrosion behaviour of coated biodegradable magnesium alloys using phase field simulation and machine learning. *Comput. Mater. Sci.* **2025**, *244*, 113546. [[CrossRef](#)]
59. Bing, H.; Liu, Z.; Ge, S.; Zou, Y.; Li, S. Phase-field modeling of stress-accelerated corrosion in Zn-Al-Mg coatings on high-strength steel wires. *Constr. Build. Mater.* **2025**, *498*, 143978. [[CrossRef](#)]
60. Lu, L.; Pestourie, R.; Yao, W.; Wang, Z.; Verdugo, F.; Johnson, S.G. Physics-informed neural networks with hard constraints for inverse design. *SIAM J. Sci. Comput.* **2021**, *43*, B1105–B1132. [[CrossRef](#)]
61. Sukumar, N.; Srivastava, A. Exact imposition of boundary conditions with distance functions in physics-informed deep neural networks. *Comput. Methods Appl. Mech. Eng.* **2022**, *389*, 114333. [[CrossRef](#)]
62. Jagtap, A.D.; Kharazmi, E.; Karniadakis, G.E. Conservative physics-informed neural networks on discrete domains for conservation laws: Applications to forward and inverse problems. *Comput. Methods Appl. Mech. Eng.* **2020**, *365*, 113028. [[CrossRef](#)]
63. Jagtap, A.D.; Karniadakis, G.E. 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*, 2002–2041. [[CrossRef](#)]
64. Moseley, B.; Markham, A.; Nissen-Meyer, T. Finite basis physics-informed neural networks (FBPINNs): A scalable domain decomposition approach for solving differential equations. *Adv. Comput. Math.* **2023**, *49*, 62. [[CrossRef](#)]
65. Kharazmi, E.; Zhang, Z.; Karniadakis, G.E. hp-VPINNs: Variational physics-informed neural networks with domain decomposition. *Comput. Methods Appl. Mech. Eng.* **2021**, *374*, 113547. [[CrossRef](#)]
66. Berrone, S.; Canuto, C.; Pintore, M. Variational Physics Informed Neural Networks: The Role of Quadratures and Test Functions. *J. Sci. Comput.* **2022**, *92*, 100. [[CrossRef](#)]
67. Samaniego, E.; Anitescu, C.; Goswami, S.; Nguyen-Thanh, V.M.; Guo, H.; Hamdia, K.; Zhuang, X.; Rabczuk, T. An energy approach to the solution of partial differential equations in computational mechanics via machine learning: Concepts, implementation and applications. *Comput. Methods Appl. Mech. Eng.* **2020**, *362*, 112790. [[CrossRef](#)]
68. E, W.; Yu, B. The Deep Ritz Method: A deep learning-based numerical algorithm for solving variational problems. *Commun. Math. Stat.* **2018**, *6*, 1–12. [[CrossRef](#)]
69. Mitusch, S.K.; Funke, S.W.; Kuchta, M. Hybrid FEM–NN models: Combining artificial neural networks with the finite element method. *J. Comput. Phys.* **2021**, *446*, 110651. [[CrossRef](#)]
70. Brunton, S.L.; Proctor, J.L.; Kutz, J.N. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proc. Natl. Acad. Sci. USA* **2016**, *113*, 3932–3937. [[CrossRef](#)] [[PubMed](#)]
71. Tartakovsky, A.M.; Ortiz Marrero, C.; Perdikaris, P.; Tartakovsky, G.D.; Barajas-Solano, D. Physics-Informed Deep Neural Networks for Learning Parameters and Constitutive Relationships in Subsurface Flow Problems. *Water Resour. Res.* **2020**, *56*, e2019WR026731. [[CrossRef](#)]
72. Song, Y.; Wang, H.; Yang, H.; Taccari, M.L.; Chen, X. Loss-attentional physics-informed neural networks. *J. Comput. Phys.* **2024**, *501*, 112781. [[CrossRef](#)]
73. Lin, S.; Chen, Y. Causality-guided adaptive sampling method for physics-informed neural networks solving forward problems of partial differential equations. *Phys. D Nonlinear Phenom.* **2025**, *481*, 134878. [[CrossRef](#)]
74. Lu, L.; Jin, P.; Pang, G.; Zhang, Z.; Karniadakis, G.E. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nat. Mach. Intell.* **2021**, *3*, 218–229. [[CrossRef](#)]
75. Li, Z.; Kovachki, N.; Azizzadenesheli, K.; Liu, B.; Bhattacharya, K.; Stuart, A.; Anandkumar, A. Fourier neural operator for parametric partial differential equations. *arXiv* **2021**, arXiv:2010.08895. [[CrossRef](#)]
76. Kovachki, N.; Li, Z.; Liu, B.; Azizzadenesheli, K.; Bhattacharya, K.; Stuart, A.; Anandkumar, A. Neural operator: Learning maps between function spaces with applications to PDEs. *J. Mach. Learn. Res.* **2023**, *24*, 1–97.
77. Li, Z.; Zheng, H.; Kovachki, N.; Jin, D.; Chen, H.; Liu, B.; Azizzadenesheli, K.; Anandkumar, A. Physics-Informed Neural Operator for Learning Partial Differential Equations. *ACM/IMS J. Data Sci.* **2024**, *1*, 1–27. [[CrossRef](#)]
78. Bronstein, M.M.; Bruna, J.; Cohen, T.; Velickovic, P. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv* **2021**, arXiv:2104.13478. [[CrossRef](#)]
79. Satorras, V.G.; Hoogeboom, E.; Welling, M. E(n) equivariant graph neural networks. *arXiv* **2021**, arXiv:2102.09844. [[CrossRef](#)]
80. Pfaff, T.; Fortunato, M.; Sanchez-Gonzalez, A.; Battaglia, P.W. Learning mesh-based simulation with graph networks. *arXiv* **2021**, arXiv:2010.03409. [[CrossRef](#)]
81. Brandstetter, J.; Worrall, D.; Welling, M. Message passing neural PDE solvers. *arXiv* **2022**, arXiv:2202.03376. [[CrossRef](#)]
82. Greydanus, S.; Dzamba, M.; Yosinski, J. Hamiltonian neural networks. *Adv. Neural Inf. Process. Syst.* **2019**, *32*, 15379–15389.
83. Jin, P.; Zhang, Z.; Zhu, A.; Tang, Y.; Karniadakis, G.E. SympNets: Intrinsic structure-preserving symplectic networks for identifying Hamiltonian systems. *Neural Netw.* **2020**, *132*, 166–179. [[CrossRef](#)] [[PubMed](#)]
84. Abdar, M.; Pourpanah, F.; Hussain, S.; Rezazadegan, D.; Liu, L.; Ghavamzadeh, M.; Fieguth, P.; Cao, X.; Khosravi, A.; Acharya, U.R.; et al. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Inf. Fusion* **2021**, *76*, 243–297. [[CrossRef](#)]

85. Psaros, A.F.; Meng, X.; Zou, Z.; Guo, L.; Karniadakis, G.E. Uncertainty quantification in scientific machine learning: Methods, metrics, and comparisons. *J. Comput. Phys.* **2023**, *477*, 111902. [\[CrossRef\]](#)
86. Yang, L.; Meng, X.; Karniadakis, G.E. B-PINNs: Bayesian physics-informed neural networks for forward and inverse PDE problems with noisy data. *J. Comput. Phys.* **2021**, *425*, 109913. [\[CrossRef\]](#)
87. Yu, X.; Wang, Q.; Khodadadi, N.; Nanni, A.; Mosalam, K.M. An adaptive Bayesian Physics-Informed neural network model for predicting the fatigue life of high-strength steel under non-uniform corrosion. *Comput.-Aided Civ. Infrastruct. Eng.* **2026**, *50*, 100112. [\[CrossRef\]](#)
88. Gal, Y.; Ghahramani, Z. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. *arXiv* **2016**, arXiv:1506.02142. [\[CrossRef\]](#)
89. Lakshminarayanan, B.; Pritzel, A.; Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. *Adv. Neural Inf. Process. Syst.* **2017**, *30*, 6402–6413.
90. Perdikaris, P.; Raissi, M.; Damianou, A.; Lawrence, N.D.; Karniadakis, G.E. Nonlinear information fusion algorithms for data-efficient multi-fidelity modelling. *Proc. R. Soc. A Math. Phys. Eng. Sci.* **2017**, *473*, 20160751. [\[CrossRef\]](#) [\[PubMed\]](#)
91. Meng, X.; Karniadakis, G.E. A composite neural network that learns from multi-fidelity data: Application to function approximation and inverse PDE problems. *J. Comput. Phys.* **2020**, *401*, 109020. [\[CrossRef\]](#)
92. Kennedy, M.C.; O'Hagan, A. Bayesian calibration of computer models. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **2001**, *63*, 425–464. [\[CrossRef\]](#)
93. Huang, T.; Li, X.; Zhang, Y.; Yao, L.; Zhang, T. Multifactorial prediction of corrosion fatigue crack growth in aluminum alloys using physics-informed neural networks. *Eng. Fail. Anal.* **2025**, *174*, 109521. [\[CrossRef\]](#)
94. Sheikh, M.F.; Kamal, K.; Rafique, F.; Sabir, S.; Zaheer, H.; Khan, K. Corrosion detection and severity level prediction using acoustic emission and machine learning based approach. *Ain Shams Eng. J.* **2021**, *12*, 3891–3903. [\[CrossRef\]](#)
95. RajKumar, A.; Paluru, N.; Mathew, R.S.; Shenai, P.; Abdeen, D.; Laycock, N.; Yalavarthy, P.K. Unsupervised machine learning for automated corrosion staging using optical microscopy images. *npj Mater. Degrad.* **2025**, *9*, 83. [\[CrossRef\]](#)
96. Vangrunderbeek, V.; Coelho, L.B.; Meeusen, M.; Terryn, H.; Mamme, M.H. A field-based computer vision and corrosion sensor approach to study electrolyte dynamics in atmospheric corrosion. *npj Mater. Degrad.* **2025**, *10*, 3. [\[CrossRef\]](#)
97. Chen, W.; Ma, L.; Li, Y.; Wu, D.; Zhou, K.; Wang, J.; Chen, Z.; Guo, X.; Li, Z.; Chowwanonthapunya, T.; et al. Prediction of coating degradation based on Environmental Factors–Physical Property–Corrosion Failure two-stage machine learning. *npj Mater. Degrad.* **2025**, *9*, 67. [\[CrossRef\]](#)
98. Samanta, S.; Luigjes, B.; Lyon, S.; Kell, D.B. Novel corrosion protective coating formulations suggested by machine learning of historic test data. *npj Mater. Degrad.* **2026**, *10*, 63. [\[CrossRef\]](#)
99. Chen, N.; Cui, C.; Ma, R.; Chen, A.; Wang, S. Sharp-Pinns: Staggered Hard-Constrained Physics-Informed Neural Networks for Phase Field Modelling of Corrosion. *Comput. Methods Appl. Mech. Eng.* **2025**, *447*, 118346. [\[CrossRef\]](#)
100. Khayatadz, M.; De Pue, L.; De Waele, W. Detection of corrosion on steel structures using automated image processing. *Dev. Built Environ.* **2020**, *3*, 100022. [\[CrossRef\]](#)
101. Rossouw, E.; Doorsamy, W. Predictive Maintenance Framework for Cathodic Protection Systems Using Data Analytics. *Energies* **2021**, *14*, 5805. [\[CrossRef\]](#)
102. Bouzaffour, K.; Lescop, B.; Gallée, F.; Talbot, P.; Rioual, S. Development of Autonomous UHF RFID Sensors Embedded in Concrete for the Monitoring of Infrastructures in Marine Environments. In *European Workshop on Structural Health Monitoring; Lecture Notes in Civil Engineering*; Springer International Publishing: Cham, Switzerland, 2021; pp. 447–454. [\[CrossRef\]](#)
103. Tuninetti, V.; Huentemilla, M.; Gómez, Á.; Oñate, A.; Menacer, B.; Narayan, S.; Montalba, C. Evaluating Pipeline Inspection Technologies for Enhanced Corrosion Detection in Mining Water Transport Systems. *Appl. Sci.* **2025**, *15*, 1316. [\[CrossRef\]](#)
104. Cheng, Y.; Wang, L.; Yang, C.; Bai, Y.; Wang, H.; Cheng, W.; Tiyyagura, H.R.; Komissarov, A.; Shin, K.S. A brief review of machine learning-assisted Mg alloy design, processing, and property predictions. *J. Mater. Res. Technol.* **2024**, *30*, 8108–8127. [\[CrossRef\]](#)
105. Toor, I.U. Artificial intelligence and machine learning in corrosion inhibitor design & development: Advances, challenges, and future perspectives. *Int. J. Electrochem. Sci.* **2026**, *21*, 101274. [\[CrossRef\]](#)
106. Xie, T.; Grossman, J.C. Crystal Graph Convolutional Neural Networks for an Accurate and Interpretable Prediction of Material Properties. *Phys. Rev. Lett.* **2018**, *120*, 145301. [\[CrossRef\]](#) [\[PubMed\]](#)
107. Reiser, P.; Neubert, M.; Eberhard, A.; Torresi, L.; Zhou, C.; Shao, C.; Metni, H.; van Hoesel, C.; Schopmans, H.; Sommer, T.; et al. Graph neural networks for materials science and chemistry. *Commun. Mater.* **2022**, *3*, 93. [\[CrossRef\]](#) [\[PubMed\]](#)
108. Amirudin, A.; Thierry, D. Application of electrochemical impedance spectroscopy to study the degradation of polymer-coated metals. *Prog. Org. Coat.* **1995**, *26*, 1–28. [\[CrossRef\]](#)
109. ISO 9223:2012; Corrosion of Metals and Alloys—Corrosivity of Atmospheres—Classification, Determination and Estimation. ISO: Geneva, Switzerland, 2012.
110. ISO 12944-2:2017; Paints and Varnishes—Corrosion Protection of Steel Structures by Protective Paint Systems—Part 2: Classification of Environments. ISO: Geneva, Switzerland, 2017.

111. Atha, D.J.; Jahanshahi, M.R. Evaluation of deep learning approaches based on convolutional neural networks for corrosion detection. *Struct. Health Monit.* **2018**, *17*, 1110–1128. [[CrossRef](#)]
112. Li, X.; Zhang, D.; Liu, Z.; Li, Z.; Du, C.; Dong, C. Share corrosion data. *Nature* **2015**, *527*, 441–442. [[CrossRef](#)] [[PubMed](#)]
113. Edelsbrunner, H.; Letscher, D.; Zomorodian, A. Topological Persistence and Simplification. *Discret. Comput. Geom.* **2002**, *28*, 511–533. [[CrossRef](#)]
114. Ghrist, R. Barcodes: The persistent topology of data. *Bull. Am. Math. Soc.* **2008**, *45*, 61–75. [[CrossRef](#)]
115. Clough, J.R.; Byrne, N.; Oksuz, I.; Zimmer, V.A.; Schnabel, J.A.; King, A.P. A Topological Loss Function for Deep-Learning Based Image Segmentation Using Persistent Homology. *IEEE Trans. Pattern Anal. Mach. Intell.* **2022**, *44*, 8766–8778. [[CrossRef](#)] [[PubMed](#)]
116. Hernández, Q.; Badías, A.; González, D.; Chinesta, F.; Cueto, E. Structure-preserving neural networks. *J. Comput. Phys.* **2021**, *426*, 109950. [[CrossRef](#)]
117. Hernández, Q.; Badías, A.; Chinesta, F.; Cueto, E. Thermodynamics-Informed Graph Neural Networks. *IEEE Trans. Artif. Intell.* **2024**, *5*, 967–976. [[CrossRef](#)]
118. Zhang, Z.; Shin, Y.; Karniadakis, G.E. GFNNs: GENERIC formalism informed neural networks for deterministic and stochastic dynamical systems. *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.* **2022**, *380*, 20210207. [[CrossRef](#)] [[PubMed](#)]
119. Park, J.S.R.; Cheung, S.W.; Choi, Y.; Shin, Y. tLaSDI: Thermodynamics-informed latent space dynamics identification. *Comput. Methods Appl. Mech. Eng.* **2024**, *429*, 117144. [[CrossRef](#)]
120. Morrison, P.J. A paradigm for joined Hamiltonian and dissipative systems. *Phys. D Nonlinear Phenom.* **1986**, *18*, 410–419. [[CrossRef](#)]
121. Grmela, M.; Öttinger, H.C. Dynamics and thermodynamics of complex fluids. I. Development of a general formalism. *Phys. Rev. E* **1997**, *56*, 6620–6632. [[CrossRef](#)]
122. Schütt, K.T.; Sauceda, H.E.; Kindermans, P.J.; Tkatchenko, A.; Müller, K.R. SchNet: A deep learning architecture for molecules and materials. *J. Chem. Phys.* **2018**, *148*, 241722. [[CrossRef](#)] [[PubMed](#)]
123. Chen, C.; Ye, W.; Zuo, Y.; Zheng, C.; Ong, S.P. Graph Networks as a Universal Machine Learning Framework for Molecules and Crystals. *Chem. Mater.* **2019**, *31*, 3564–3572. [[CrossRef](#)]
124. Batzner, S.; Musaelian, A.; Sun, L.; Geiger, M.; Mailoa, J.P.; Kornbluth, M.; Molinari, N.; Smidt, T.E.; Kozinsky, B. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nat. Commun.* **2022**, *13*, 2453. [[CrossRef](#)] [[PubMed](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.