NETZERO 10.66967/netzero.2026.v2i105 CoSci Press

Research Article

Probabilistic Digital Twin Framework for Climate-Resilient Offshore Infrastructure Under Coupled Corrosion–Fatigue Degradation and Environmental Uncertainty

Seyed Reza Samaei [1,*], Mohammad Asadian Ghahfarokhi [1], Nader Bavanpouri [2], Mohamad Amein Amjadi Kelishemi [1], Sajjad Pendashteh [1]

[1] Department of Marine Industries, Science and Research Branch (SR.C.), Islamic Azad University, Tehran 1477893855, Iran.
[2] Department of Maritime Business Management, Science and Research Branch (SR.C.), Islamic Azad University, Tehran 1477893855, Iran.
* Corresponding author: samaei@srbiau.ac.ir

Article History: Received 29 May 2026; Revised 31 July 2026; Accepted 11 October 2026; Published 16 October 2026

Abstract: Offshore assets deteriorate through mechanisms that reinforce one another: corrosion causes section loss and raises the local stress, the elevated stress accelerates fatigue, and fatigue damage exposes fresh metal to the marine environment. Deterministic integrity assessment cannot represent that feedback or the environmental variability that drives it. This paper presents a physics-informed probabilistic digital twin for offshore infrastructure in which coupled corrosion–fatigue degradation modelling, Monte Carlo uncertainty propagation, sequential Bayesian data assimilation and a physics-constrained neural surrogate are combined within a single lifecycle assessment loop. The degradation state is a bounded composite index of corrosion penetration, Miner fatigue damage, cumulative environmental exposure and an interaction component, with the stress-amplification coupling regularised so that no governing relation becomes singular as failure is approached. Structural performance is tracked through the second-moment reliability index, and the Infrastructure Degradation Instability Index (IDII) is defined as a bounded, dimensionless ratio between an instability driver, the product of normalised degradation rate and dispersion, and the normalised reliability margin. For a 32 mm-wall tubular reference component configured deliberately as an aggressive 30-year demonstration case, the reliability index falls from 3.48 to 0.70 and the target of 1.5 is reached at 17.8 years; the empirical failure probability counted from the 30,000 Monte Carlo limit-state samples differs from $\Phi(-\beta)$ by at most 2.2×10^{-4}, so the two-moment approximation is verified rather than assumed. Global variance decomposition attributes 87% of the variance in end-of-life degradation to environmental severity. Raising the environment-driven variability multiplier from 0.5 to 2.0 at constant mean severity advances the reliability-target year from 21.5 to 10.8 even though the mean degradation trajectory falls, so variability acts through dispersion rather than through the mean. Sequential importance resampling of annual thickness and strain measurements reduces the predictive dispersion of the 30-year degradation state by 81.5%. IDII thresholds of 0.15, 0.35, 0.50 and 0.65 are mapped to reliability-index levels of 3.06, 2.26, 1.71 and 1.05, and a 21-case sensitivity study shows that the composite-index weights move the reliability-target year between 15.9 and 19.3 years while leaving the regime ordering and the warning lead time (6.6–8.1 years) essentially unchanged. The surrogate reproduces held-out trajectories with RMSE 0.0078 and R^2 0.999 but is unreliable beyond its training horizon, which is why long-horizon inference is carried by the Bayesian filter. All parameters, distributions and hyperparameters are reported; the framework is presented as probabilistic decision support for a demonstration configuration, with field validation identified as the prerequisite for operational use.

Keywords: digital twin; offshore infrastructure; corrosion–fatigue interaction; Bayesian data assimilation; structural reliability; climate resilience

1. Introduction

Fixed and floating offshore structures underpin a large share of global energy supply and maritime logistics, and an increasing proportion of them are operating beyond their original design intent. Ageing jackets, monopiles and platform legs are being life-extended at the same time as the metocean climate they were designed against is shifting. The combination places unusual demands on integrity management: operators must decide when to inspect, repair or decommission assets whose deterioration is driven by processes that interact, and whose loading environment is itself non-stationary [1–3].

Conventional integrity assessment answers those questions deterministically. Environmental loading, corrosion allowance, fatigue capacity and operational duty are represented by characteristic values, safety factors absorb the residual uncertainty, and the resulting remaining-life estimate is a single number. The approach has served the industry well for design, but it has two structural weaknesses when applied to in-service decision-making. First, it cannot propagate the dispersion that dominates late-life behaviour, so it gives no basis for judging how much confidence a remaining-life estimate deserves. Second, because each degradation mechanism is assessed on its own, the interaction between them is absorbed into conservatism rather than modelled [4,5].

That interaction is not a second-order effect. Corrosion reduces the load-bearing section and sharpens the local stress field, which shortens fatigue life; fatigue cracking and coating damage expose fresh metal, which accelerates corrosion [6–8]. Treating the two mechanisms additively understates the rate at which reliability is consumed in the final third of service life, precisely the period in which life-extension decisions are made.

Digital twin technology offers a route past both weaknesses. Since the concept was articulated for product lifecycle management [9] and for aerospace vehicle certification [10], it has developed into a framework in which a numerical model is kept synchronised with an asset through monitoring data [11,12]. Offshore applications now span floating wind turbine damage detection, platform monitoring, crane integrity and vessel lifecycle management [13–16]. Yet most reported implementations remain deterministic and visualisation-oriented. They display condition;

they do not quantify how uncertain that condition is, nor how quickly the uncertainty is growing.

Three specific gaps motivate the present work. Uncertainty is rarely propagated end-to-end, so digital twin outputs cannot support risk-based decisions [17,18]. Coupled corrosion–fatigue physics is rarely embedded in the twin itself, so the model that is being updated is not the model that governs the failure mode. And machine learning components are frequently described in general terms—"AI-assisted prediction"—without the architecture, training data or validation evidence that would let a reader judge whether the component is doing useful work, or whether it is extrapolating beyond any regime in which it was trained [19].

There is also a measurement problem that has received less attention than it deserves. Reliability metrics such as $\beta(t)$ describe how close a structure is to a limit state, but they say nothing about how fast that margin is being consumed or how confidently the estimate is known. Two assets with identical reliability indices can be in very different situations if one is degrading slowly with tight uncertainty bounds and the other is degrading rapidly with wide ones. Capturing that distinction requires an indicator that combines rate, dispersion and margin.

This paper develops such a framework and reports it in full. The specific contributions are:

1. **A coupled corrosion–fatigue degradation model with explicit, bidirectional interaction terms** and a regularised stress-amplification law that remains bounded as the section loss approaches the wall thickness, removing the numerical singularity present in conventional formulations.
2. **An Infrastructure Degradation Instability Index** defined as a bounded dimensionless ratio, together with engineering thresholds calibrated against reliability-index levels rather than asserted qualitatively, a 21-case sensitivity analysis of the composite-index weights, and a demonstration of its sensitivity to environmental variability.
3. **A fully specified Bayesian data assimilation scheme**—priors, likelihood, sequential importance-resampling algorithm, regularisation and update frequency—benchmarked against a non-updating baseline across 150 synthetic assets rather than a single illustrative case.
4. **A physics-constrained neural surrogate reported with complete architecture, hyperparameters and error metrics**, including an explicit extrapolation study that establishes where the surrogate may and may not be trusted, and a corresponding statement of its bounded role in the architecture.
5. **A global variance decomposition** identifying which uncertainty sources actually govern end-of-life degradation and reliability, an independent empirical verification of the two-moment failure probability against the Monte Carlo limit-state samples, and a complete parameter and distribution specification enabling independent reproduction.

The paper is organised as follows. Section 2 reviews digital twin, corrosion–fatigue, probabilistic assessment and physics-informed machine learning literature, and states the research gap. Section 3 develops the framework and its governing relations. Section 4 specifies the reference component, the environmental model and the simulation procedure. Section 5 details data assimilation and the surrogate. Section 6 reports results. Section 7 discusses implications, scalability and limitations. Section 8 concludes.

2. Literature Review and Research Gap

2.1. Digital Twins for Offshore Structural Integrity

The digital twin concept originated in product lifecycle management, where Grieves and Vickers (2017) framed it as a means of anticipating emergent behaviour in complex systems [9], and was given engineering substance by Glaessgen and Stargel (2012), who proposed a vehicle-specific, continuously updated simulation as the basis for certification of future NASA and US Air Force airframes [10]. Subsequent work generalised the architecture to manufacturing and service systems [11,20], and Kapteyn, Pretorius and Willcox (2021) supplied a probabilistic graphical formulation that makes the asset state [12], the observations and the model parameters explicit random variables—the formulation on which uncertainty-aware twins now rest.

Offshore adoption has been rapid. Li and Brennan (2024) reviewed digital-twin-enabled structural integrity management and identified the absence of integrated uncertainty treatment as the principal barrier to risk-informed use [21]; the companion study showed that virtually monitored data can materially improve inspection planning [22]. Mousavi et al. (2024) combined deep learning with a floating wind turbine twin for damage detection [13], Pastor-Sanchez et al. (2025) demonstrated real-time structural health monitoring for floating turbines [14], and Rølvåg and Stranden (2022) applied the approach to offshore crane structures [15]. Mauro and Kana (2023) surveyed maritime applications and reached a similar conclusion to Li and Brennan: most twins report condition rather than risk [16]. Sakr and Sadhu (2024) reached the same conclusion for civil infrastructure more broadly [23], and Lai et al. (2023) showed how measured and computed data can be fused within a twin to improve state estimation [24]. Chen et al. (2023) surveyed the specific role of machine learning in predictive-maintenance twins and noted that its contribution is rarely quantified [25]. Related offshore monitoring studies have demonstrated high-density sensing and machine-learning damage detection for wharves, jack-up legs and turbine support structures [26–29], and adaptive real-time digital-twin monitoring of offshore marine structures [30].

2.2. Corrosion–Fatigue Interaction

Marine corrosion of structural steel is conventionally described by a power law in exposure time, with Melchers (2003) providing the phenomenological model and the accompanying uncertainty characterisation that remain the reference for immersion conditions [31,32]. Fatigue assessment relies on the Palmgren–Miner accumulation rule applied to S–N curves [33], or on fracture-mechanics propagation from the Paris–Erdogan law [34].

The interaction between the two has been recognised for decades but is still not routinely modelled. Adedipe, Brennan and Kolios (2016) reviewed corrosion fatigue in offshore structures and set out the challenges specific to the offshore wind sector [6]; Larrosa, Akid and Ainsworth (2018) surveyed damage-tolerance models for corrosion fatigue and noted that most treat the environment as a multiplier on crack growth rather than as a coupled state [8]. Shamir, Braithwaite and Mehmanparast (2023) quantified the fatigue life reduction caused by corrosion pits in offshore wind support structures [7], and Shittu et al. (2020, 2021) developed probabilistic damage-tolerance formulations for pitting corrosion–fatigue, comparing S–N and fracture-mechanics routes [4,35]. Tafazzoli Moghaddam et al. (2020), Zhu et al. (2024), Yeter and Brennan (2024) and Partovi-Mehr et al. (2024) extended reliability assessment to floating substructures and to measured offshore duty at the Block Island Wind Farm [36–39]. What these studies share is a treatment of corrosion and fatigue as sequential rather than mutually reinforcing: corrosion sets a thickness, fatigue is then evaluated on that thickness. The reverse path—fatigue damage accelerating corrosion—is seldom closed.

2.3. Probabilistic Assessment and Bayesian Updating

Probabilistic structural assessment replaces characteristic values with distributions and reports a reliability index or failure probability. Beck and Katafygiotis (1998) established the Bayesian statistical framework for updating structural models and their uncertainties [40], and Katafygiotis and Beck (1998) analysed the identifiability conditions under which such updating is well posed [41]; Vanik, Beck and Au (2000) applied the framework to structural health monitoring [42]. These papers define the inferential problem that modern digital twins solve numerically.

Computationally, posterior evaluation for nonlinear, non-Gaussian degradation models is normally handled by sequential Monte Carlo. Arulampalam et al. (2002) give the standard treatment of particle filtering, including the degeneracy problem and the resampling strategies that address it [43]. Rocchetta et al. (2018) demonstrated online Bayesian model updating for structural health monitoring [44], and Torzoni, Manzoni and Mariani (2025) improved updating efficiency through learnable mappings [45]. In offshore contexts, Wang et al. (2021) reviewed reliability analysis of turbine support structures under digital twin solutions [46], and Jorgensen et al. (2023) set out the requirements for applying the paradigm to uncertain fatigue analysis [17].

Two shortcomings recur. Updating is often applied to a single scalar observable rather than to the joint parameter set of a coupled degradation model, and the computational method—Kalman variant, particle filter, or Markov chain Monte Carlo—is frequently unstated, which prevents assessment of whether the posterior is credible.

2.4. Physics-Informed Machine Learning

Purely data-driven surrogates interpolate well and extrapolate badly, a property that matters acutely for degradation forecasting because the regime of interest is by definition one that has not yet been observed. Physics-informed machine learning addresses this by embedding governing relations into the learning problem, either as loss penalties or as architectural constraints. Raissi, Perdikaris and Karniadakis (2019) introduced physics-informed neural networks for forward and inverse problems governed by partial differential equations [47], and Karniadakis et al. (2021) surveyed the field and its failure modes [19]. Viana and Subramaniyan (2021) reviewed the relationship between Bayesian calibration and physics-informed networks in scientific modelling, which is directly relevant to hybrid twins that do both [48].

For degradation specifically, Dourado and Viana (2020) used a physics-informed network to estimate *missing physics* in cumulative damage models, with corrosion fatigue as the case study: a known cumulative-damage kernel is retained and the network learns only the unmodelled correction [49]. That decomposition—known physics plus learned residual—is the structure adopted in Section 5.3 of the present work. In offshore energy applications, Zhang et al. (2024) combined a load surrogate model with a genetic algorithm for offshore wind farm layout optimisation, illustrating the computational leverage that surrogates provide when the underlying physical evaluation is expensive [50].

The offshore literature has, however, been slow to report the diagnostics that make such components assessable. Architecture, training set size, hyperparameters, held-out error metrics and—most importantly—extrapolation behaviour are frequently omitted, so it is impossible to tell whether a physics-informed label reflects a genuine constraint or a nominal one.

2.5. Research Gap

The review supports five statements. Offshore digital twins are predominantly deterministic and report condition rather than risk. Corrosion–fatigue coupling is usually unidirectional where it is modelled at all. Bayesian updating is often described without specifying priors, likelihood or the posterior computation method. Machine learning components are reported without the evidence needed to bound their validity. And no established indicator combines degradation rate, degradation uncertainty and reliability margin into a single quantity with calibrated engineering thresholds. The framework developed below addresses each of these in turn, and reports the numerical detail required for independent verification.

3. Physics-Informed Probabilistic Digital Twin Framework

3.1. Architecture

The framework is organised as five interacting layers. A sensing layer acquires structural and environmental observations. A multi-physics layer evolves the coupled corrosion–fatigue degradation state. A probabilistic layer propagates parameter uncertainty through that model by Monte Carlo simulation. An assimilation layer conditions the model parameters on incoming measurements by sequential Bayesian inference. A decision layer converts the resulting posterior state into reliability, instability and inspection-timing outputs. A neural surrogate sits alongside the multi-physics layer and is used to accelerate repeated evaluation within the envelope in which it was trained.

What distinguishes this arrangement from a monitoring dashboard is that every layer carries a distribution rather than a value, and that the object being updated by the assimilation layer is the parameter set of the physical degradation model—not a curve fitted to the data. The asset-specific information contained in inspection records therefore propagates into the forward projection through the physics, which is what allows the twin to forecast beyond the period in which it has been observed.

Figure 1 shows the architecture and the flow of information between layers.

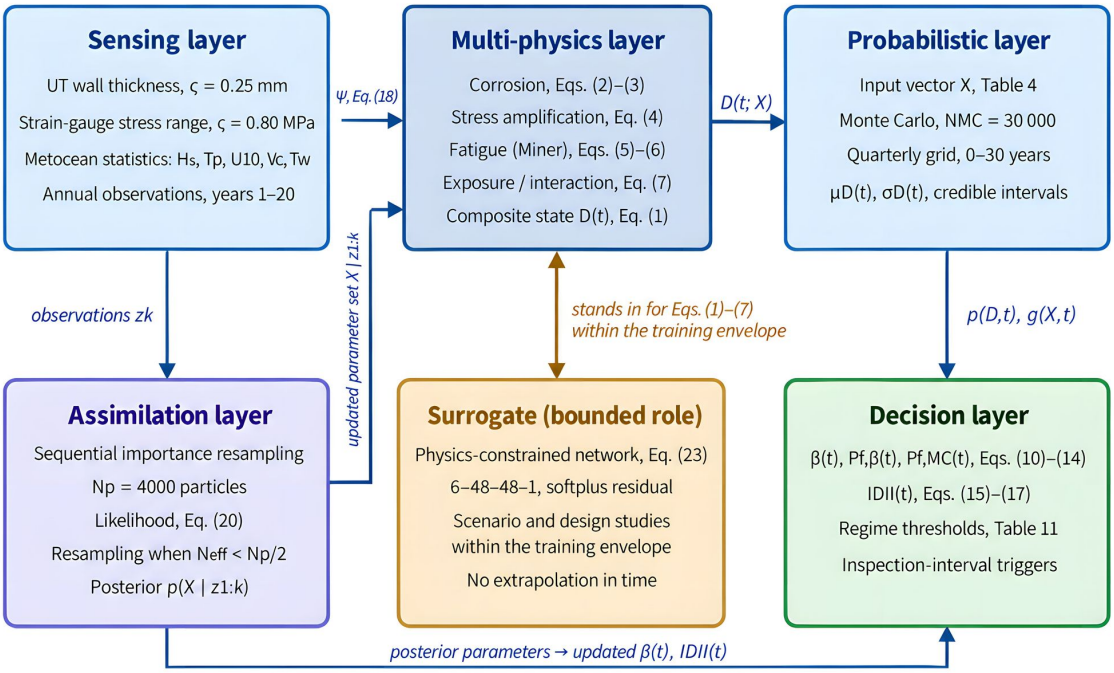


Figure 1. Architecture of the physics-informed probabilistic digital twin as implemented in this study. Each block lists only the quantities, equations and tables used in the analysis; arrows show the direction of information flow: observations enter the assimilation layer, the posterior parameter set is returned to the physics model and passed to the decision layer, and the surrogate stands in for the multi-physics model for repeated evaluation within its training envelope.

3.2. Coupled Multi-physics Degradation Model

The degradation state is expressed as a bounded composite index formed from four normalised, dimensionless components:

$$D(t) = w_c D_c^*(t) + w_f D_f^*(t) + w_e D_e^*(t) + w_i D_i^*(t), \qquad \sum_j w_j = 1, \quad 0 \le D \le 1 \tag{1}$$

where D_c^*, D_f^*, D_e^* and D_i^* denote normalised corrosion, fatigue, environmental-exposure and interaction damage. Each component is defined explicitly below, while the baseline weighting scheme and its parameterisation are presented in Section 4.4. The influence of the component weights on the results is quantified in Section 6.7. Throughout this paper $D(t)$ is referred to as the *degradation state*, and the term *instability index* is reserved exclusively for $IDII(t)$ defined in Section 3.5. The two quantities are distinct: D measures how much damage has accumulated, *IDII* measures how unstably that damage is evolving relative to the remaining reliability margin.

Corrosion. Penetration depth follows a power law in exposure time, in the form established for marine immersion by Melchers (2003), modified by an initiation delay representing coating and cathodic-protection life and by a feedback term through which accumulated fatigue damage accelerates metal loss [31]:

$$d_c(t) = c_\alpha\, \alpha_c\, \Psi^{1/2} \langle t - t_i \rangle^{n_c} \left[1 + a_{fc} D_f^*(t)\right] \tag{2}$$

where $\langle\cdot\rangle$ denotes the positive part, Ψ is the normalised environmental severity index defined in Section 4.3, a_{fc} is the fatigue-to-corrosion coupling coefficient acting on the capped Miner damage of Equation (6), and c_α is a dimensionless calibration multiplier on the penetration depth (Section 4.4). Normalisation is by the corrosion allowance t_{ca} rather than the full wall thickness, so that $D_c^* = 1$ corresponds to exhaustion of the design allowance:

$$D_c^*(t) = min\left(\frac{d_c(t)}{t_{ca}},\, 1\right) \tag{3}$$

Stress amplification. Section loss raises the local stress. Writing the section-loss fraction as $\theta(t) = d_c(t)/t_0$, the elementary thin-wall result gives $K_\sigma = (1-\theta)^{-1}$, which diverges as $\theta \to 1$. Because this singularity propagates into the fatigue rate and hence into the reliability estimate, it is replaced by a bounded rational form with the same small-θ behaviour:

$$K_\sigma(t) = 1 + \left(K_{\sigma,max} - 1\right)\frac{\theta(t)}{\theta(t) + c_K} \tag{4}$$

Equation (4) is monotonic, satisfies $K_\sigma(0) = 1$ and saturates at $K_{\sigma,max}$; with $c_K = 0.08$ it rises more steeply than the thin-wall expression at small section loss and remains finite where that expression diverges. The saturation value is the stress concentration factor adopted for the maximum admissible section loss of the detail, and c_K sets the section loss at which half of the amplification is developed, so both constants have a physical meaning rather than being numerical conveniences. Figure 2b compares the two forms.

Fatigue. Damage accumulates by the Palmgren–Miner rule applied to a single-slope S–N curve, with the effective stress range carrying both the corrosion-induced amplification and the environmental severity:

$$\Delta\sigma_{eff}(t) = c_\sigma\, K_\sigma(t)\, \Delta\sigma_0\, \Psi^{0.8} \tag{5}$$

$$\frac{dD_f}{dt} = \frac{\nu\left[\Delta\sigma_{eff}(t)\right]^m}{\bar{A}}, \qquad D_f^*(t) = min\left(\int_0^t \frac{dD_f}{d\tau}\, d\tau,\; 1\right) \tag{6}$$

In Equation (5), c_σ is the dimensionless calibration multiplier on the nominal stress range (Section 4.4). The exponent 0.8 on Ψ in Equation (5) reflects the near-linear scaling of wave-induced stress range with significant wave height; combined with the S–N slope $m = 3$ this makes fatigue damage strongly convex in environmental severity, which is the mechanism by which variability—not merely mean severity—drives degradation.

Environmental exposure and interaction. Cumulative

environmental exposure and the interaction contribution are

$$D_e^*(t) = min\left(\frac{t}{T_d}\Psi,\ 1\right), \qquad D_i^*(t) = \frac{K_\sigma(t) - 1}{K_{\sigma,max} - 1} \tag{7}$$

Equations (2), (5) and (6) close the bidirectional coupling: corrosion raises the stress range that drives fatigue, and fatigue damage raises the corrosion rate. Figure 2 shows the resulting component trajectories, the stress-amplification history and the contribution of each component to the end-of-life degradation state.

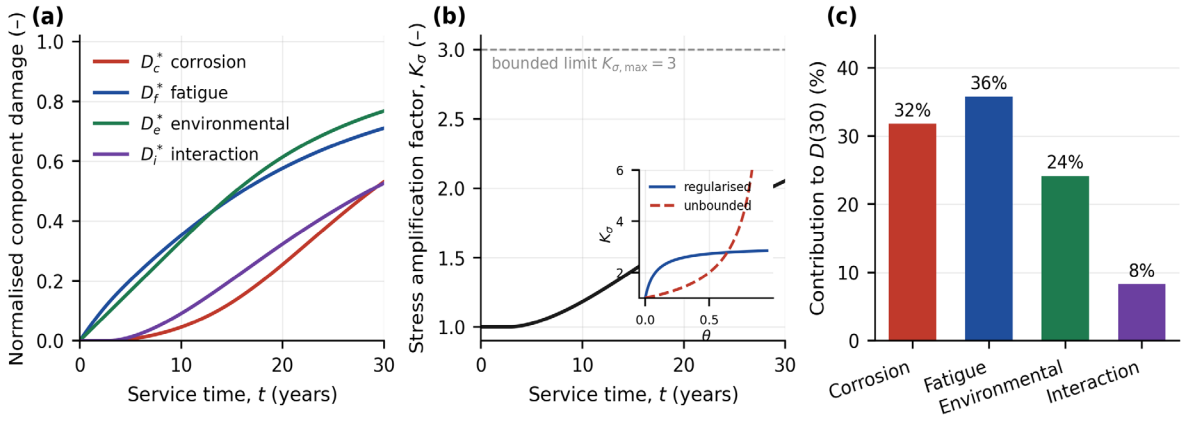


Figure 2. Coupled degradation mechanics. (**a**) Evolution of the four normalised damage components. (**b**) Stress amplification factor, with the inset comparing the regularised law of Equation (4) against the unbounded thin-wall expression. (**c**) Contribution of each component to the mean degradation state at 30 years.

3.3. Probabilistic State Representation

The uncertain input vector is $\mathbf{X} = [t_0, t_{ca}, \alpha_c, n_c, \Psi, \Delta\sigma_0, \nu, \log\bar{A}, R_0, S]^\top$, comprises the probabilistic model parameters and environmental inputs, whose marginal distributions and statistical properties are specified in Section 4.4. Uncertainty is propagated by direct Monte Carlo simulation, so that the time-dependent distribution of the degradation state is

$$p(D,t) = \int p(D \mid \mathbf{X}, t)p(\mathbf{X})\, d\mathbf{X} \tag{8}$$

evaluated with $N_{MC} = 30{,}000$ realisations on a quarterly time grid over the 30-year design life. Convergence was verified by comparing the mean and standard deviation of $D(30)$ across independent batches; both were stable to within 0.3% for $N_{MC} \geq 20{,}000$.

3.4. Reliability Formulation

The limit state compares degraded capacity against operational demand. Capacity is written as the initial capacity reduced in proportion to the degradation state through a transfer coefficient κ, which maps the composite index onto residual structural capacity:

$$g(\mathbf{X}, t) = R_0[1 - \kappa D(t)] - s_u S \tag{9}$$

where R_0 is the normalised initial capacity, S the normalised demand and s_u the utilisation ratio; κ is the fraction of the initial capacity lost when the composite index reaches unity, With the calibrated capacity-transfer coefficient κ = 0.387, complete exhaustion of the composite degradation index corresponds to a 38.7% reduction in the initial structural capacity. The reliability index is the second-moment (Cornell) index,

$$\beta(t) = \frac{\mu_g(t)}{\sigma_g(t)} = \frac{E[g(\mathbf{X},t)]}{\sqrt{Var[g(\mathbf{X},t)]}} \tag{10}$$

from which the failure probability and reliability follow as

$$P_f(t) = \Phi\big(-\beta(t)\big), \qquad R(t) = P[g(\mathbf{X},t) > 0] = 1 - P_f(t) \tag{11}$$

where Φ denotes the standard normal cumulative distribution function. The failure probability obtained from the second-moment reliability index is denoted hereafter by $P_{f,\beta}$ to distinguish it from the failure probability estimated directly from the Monte Carlo limit-state samples. Because the complete empirical distribution of the limit-state function is available from the simulation, an independent non-parametric estimate is also evaluated as

$$P_{f,MC}(t) = \frac{1}{N_{MC}} \sum_{j=1}^{N_{MC}} \mathbb{1}\left[g_j(t) \leq 0\right] = \frac{N_f(t)}{N_{MC}} \tag{12}$$

where $\mathbb{1}[\cdot]$ is the indicator function and $N_f(t)$ the number of realisations in which the limit state is violated at time t. The corresponding empirical reliability is $R_{MC}(t) = 1 - P_{f,MC}(t)$. The discrepancy between the two estimates is quantified by the absolute difference

$$\Delta P_f(t) = \left|P_{f,MC}(t) - P_{f,\beta}(t)\right| \tag{13}$$

and the sampling uncertainty of the empirical estimate is reported as the Wilson 95% confidence interval for a binomial proportion,

$$P_{f,MC}^{\pm} = \frac{1}{1 + z^2/N_{MC}}\left[P_{f,MC} + \frac{z^2}{2N_{MC}} \pm z\sqrt{\frac{P_{f,MC}(1 - P_{f,MC})}{N_{MC}} + \frac{z^2}{4N_{MC}^2}}\right], \quad z = 1.96 \tag{14}$$

The Wilson interval is used rather than the normal approximation because it remains valid when N_f is small, as it is in early life. This comparison matters because the Cornell index depends only on the first two moments of the limit-state distribution, whereas the empirical estimate retains any skewness or other non-Gaussian features of the simulated response. Both estimates are reported side by side in Section 6.2, and $P_{f,\beta}$ is used for the remaining results because that comparison shows it to be adequate.

3.5. The Infrastructure Degradation Instability Index

Reliability alone does not describe the *stability* of a degradation process. An asset with $\beta = 2$ that is degrading slowly, with tightly bounded uncertainty, is in a materially different condition from one with the same β that is degrading rapidly with wide predictive intervals. *IDII* is introduced to separate these cases.

Two normalised quantities are formed. The instability driver combines the degradation rate and the degradation dispersion, each referred to a stated reference value:

$$\Lambda(t) = \left(\frac{\dot{D}(t)}{\dot{D}_{ref}}\right)\left(\frac{\sigma_D(t)}{\sigma_{D,ref}}\right), \qquad \dot{D}_{ref} = \frac{D_{crit}}{T_d} \tag{15}$$

The reference rate $\dot{D}_{ref}$ is the constant rate that would

consume the critical degradation level over the design life, and $\sigma_{D,ref}$ is the dispersion accepted at end of design life. The corresponding reference values are specified in Section 4.4. The reliability margin is normalised by the target reliability index:

$$\hat{\beta}(t) = \frac{\beta(t)}{\beta_{target}} \tag{16}$$

The index is then the saturating ratio

$$IDII(t) = \frac{\Lambda(t)}{\Lambda(t) + \hat{\beta}(t)} \tag{17}$$

Equation (17) has four properties that a ratio of the form $\dot{D}\sigma_D/R$ does not possess. It is dimensionless, whereas the ratio $\dot{D}\sigma_D/R$ carries units of inverse time. It is bounded in $[0,1)$ for every state with $\beta > 0$, which holds throughout the design life of the reference component, whereas dividing by $R(t)$ or by $\beta(t)$ produces a singularity as the structure approaches failure. It is exactly zero at commissioning, since $\sigma_D(0) = 0$. And it has a direct interpretation at its midpoint: $IDII = 0.5$ is the instant at which the instability driver equals the normalised reliability margin, which is to say the moment when the rate at which uncertainty and damage are accumulating first outweighs the margin still available. Equation (17) is the single definition of *IDII* used throughout this paper, in the text, in the figures and in the tables.

Each term of Equation (17) is computed as follows. $\dot{D}(t)$ is the time derivative of the Monte Carlo mean degradation state evaluated by central differences on the quarterly grid (forward difference at $t = 0$); $\sigma_D(t)$ is the sample standard deviation of D across realisations at time t; $\beta(t)$ follows from Equation (10). No additional calibration is involved beyond the three reference constants D_{crit}, $\sigma_{D,ref}$ and β_{target}.

Figure 3 shows the resulting index for the reference component, together with the crossover between Λ and $\hat{\beta}$, and the regime bands calibrated in Section 6.6.

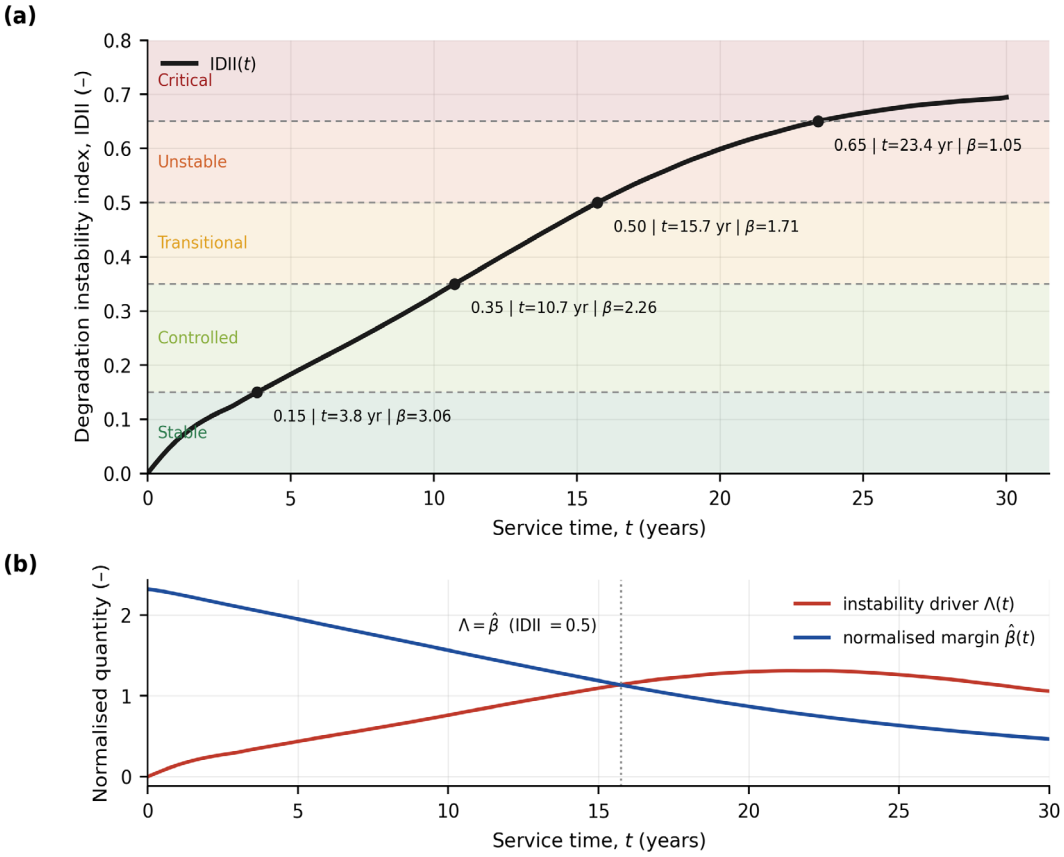


Figure 3. (**a**) Evolution of IDII for the reference component, with the five regime bands and the reliability index corresponding to each threshold crossing. (**b**) The two constituent quantities of Equation (17); their intersection defines *IDII* = 0.5.

4. Reference Component, Environmental Model and Simulation Procedure

4.1. Scope of the Case Study

The numerical study is deliberately demonstration-oriented. It uses a fully specified generic component rather than a proprietary asset, for two reasons: the parameter values can be published in full, and the behaviour of the proposed framework can be examined without confounding it with site-specific calibration. The results therefore establish the internal consistency, numerical behaviour and reproducibility of the framework for the reference configuration considered here; they do not constitute independent validation against a documented in-service offshore asset. Validation against long-term inspection and monitoring records from an operating structure remains the principal requirement for future operational deployment.

Accordingly, the decision-support outputs reported in this study should be interpreted as demonstration-case indicators rather than asset-qualified operational rules. Application to an in-service structure would require asset-specific calibration of the degradation model, independent validation against traceable field observations, verification of the applicable reliability target, and recalibration of the *IDII* thresholds and associated intervention criteria. The numerical crossing years and management transitions reported below are therefore specific to the deliberately configured reference case and are not presented as representative service-life predictions for typical code-compliant offshore assets.

4.2. Reference Component Specification

The reference component is a tubular member representative of a jacket leg or a monopile section in the submerged zone, modelled as a cantilevered cylindrical shell with a fixed base and a free head carrying the transferred topside and hydrodynamic actions; its specification is given in Table 1. Single values are used for every parameter so that the model is fully defined and reproducible.

Table 1. Geometric, material and duty specification of the reference component.

Quantity	Symbol	Value	Basis
Member length	L	22.0 m	Representative submerged bay height
Outer diameter	D_o	1.80 m	Representative jacket leg / monopile section
Initial wall thickness	t_0	32.0 mm	$D_o/t_0 \approx 56$, typical of offshore tubulars
Corrosion allowance	t_{ca}	6.0 mm	Design allowance for the submerged zone
Elastic modulus	E	210 GPa	Structural steel
Poisson's ratio	ν_p	0.30	Structural steel
Density	ρ	7850 kg m^{-3}	Structural steel
Yield strength	f_y	355 MPa	S355 grade
Boundary conditions	—	Fixed base, free head	Cantilever idealisation of a submerged bay
Design life	T_d	30 years	Assessment horizon
Utilisation ratio	s_u	0.652	Mean demand / mean initial capacity
Corrosion initiation time	t_i	3.0 years	Coating and cathodic-protection effective life
Target reliability index	β_{target}	1.5	Annual ultimate limit-state check
S–N curve	$\log_{10}\bar{A}, m$	12.164, 3.0	Single-slope curve, welded tubular detail

The parameter combination used for the reference component is intentionally selected to provide an aggressive lifecycle demonstration in which the principal reliability and IDII regime transitions occur within the 30-year numerical horizon. In particular, the utilisation, degradation and uncertainty settings are not intended to represent a population-average or nominally expected code-compliant offshore member. This deliberate configuration allows the complete decision logic of the framework to be exercised within a finite simulation period. Consequently, the baseline crossing of $\beta_{target} = 1.5$ at 17.8 years reported later is a scenario-specific demonstration result, not an estimate of the expected service life of a typical offshore asset.

4.3. Environmental Representation

Metocean conditions are represented by five variables—significant wave height H_s, peak period T_p, wind speed at 10 m U_{10}, current speed V_c and water temperature T_w—whose marginal statistics are given in Table 2. Marginals are fitted to the distribution families conventionally used for each variable, and dependence is imposed through a Gaussian copula with the correlation structure of Table 2 and Figure 4c, which reproduces the strong H_s–U_{10} association characteristic of wind-driven seas.

Table 2. Metocean variables: marginal statistics, distribution families and correlation structure.

Variable	Mean	CoV (%)	Distribution	ρ with H_s
Significant wave height, H_s (m)	2.81	56.3	Weibull	1.00
Peak period, T_p (s)	8.24	23.7	Lognormal	0.62
Wind speed, U_{10} (m s^{-1})	11.52	31.5	Weibull	0.71
Current speed, V_c (m s^{-1})	0.82	28.6	Lognormal	0.48
Water temperature, T_w (°C)	13.45	15.2	Normal	−0.21

These five variables are condensed into the single normalised environmental severity index Ψ that appears in Equations (2), (5) and (7). Ψ is a realisation-level index of the long-term severity of the metocean climate to which the member is exposed: it is drawn once per realisation, held constant within it, normalised to unit mean across the population of realisations, and follows a Weibull distribution whose shape parameter k is set by the environmental variability multiplier σ_E:

$$\Psi \sim Weibull(k), \quad \frac{\sqrt{\Gamma(1+2/k)-\Gamma^2(1+1/k)}}{\Gamma(1+1/k)} = 0.563\,\sigma_E, \quad E[\Psi] = 1 \tag{18}$$

so that the coefficient of variation of Ψ equals 0.563 σ_E exactly (k = 1.84 at σ_E = 1.0). The baseline case σ_E = 1.0 reproduces the 56.3% coefficient of variation of H_s in Table 2. Values of σ_E above unity represent populations of climates with heavier upper tails at unchanged mean severity, and are used in Section 6.3 to isolate the effect of variability from that of mean intensity. Because the corrosion-rate scatter, the hot-spot stress-range scatter and the load-model uncertainty of the reference component are all driven by the same metocean climate, Because the corrosion-rate scatter, the hot-spot stress-range scatter, and the load-model uncertainty of the reference component are driven by the same metocean climate, σ_E also scales the baseline coefficients of variation of α_c, $\Delta\sigma_0$ and S as specified in Section 4.4. The multiplier therefore represents the environment-driven variability of the problem as a whole. Within-realisation temporal variability of the climate is not resolved by the scalar Ψ; its consequences are discussed in Section 7.6.

Figure 4 summarises the environmental model.

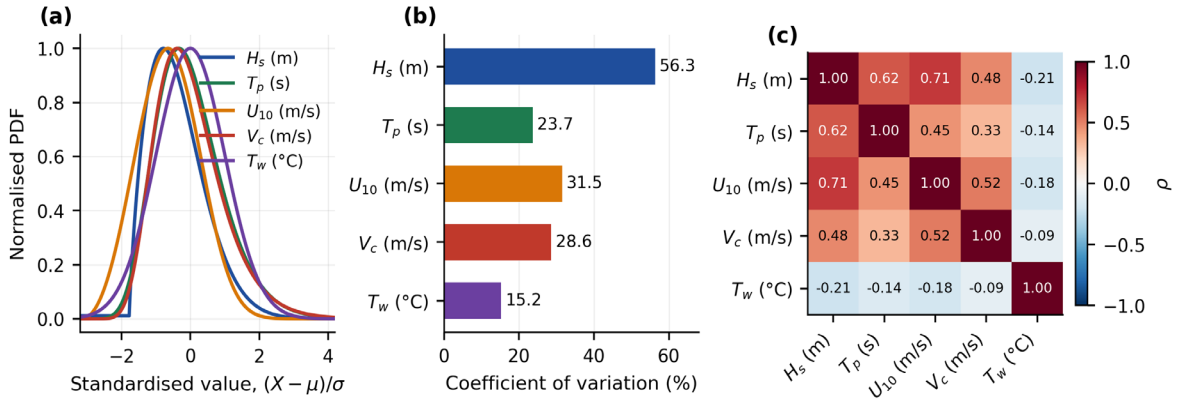


Figure 4. Environmental representation. (**a**) Standardised probability density functions (PDFs) of the five metocean variables, expressed in terms of their standardised values. (**b**) Coefficients of variation of the individual metocean variables. (**c**) Correlation matrix representing the dependence structure imposed through the Gaussian copula.

4.4. Probabilistic Simulation Procedure

Each Monte Carlo realisation samples $\mathbf{X}$, integrates Equations (2)–(7) on a quarterly grid, evaluates Equation (9), and accumulates the moments required for Equations (10)–(17). Table 3 summarises the uncertainty sources and the pathway through which each affects performance. The utilisation ratio $s_u = 0.652$ follows from Table 1 and gives $\beta(0) = 3.47$ from Equations (9) and (10) before any degradation (3.48 from the Monte Carlo simulation results reported in Section 6.2). The two calibration multipliers c_α and c_σ were determined by nonlinear least squares so that the baseline mean and standard deviation of $D(t)$ at five-year intervals reproduce the lifecycle anchors that define the aggressive demonstration case of Section 4.2 (the lifecycle mean and standard-deviation results reported in Section 6.2); the largest residual of the fit is 0.03 in μ_D and 0.025 in σ_D. The capacity transfer coefficient $\kappa = 0.387$ was then set so that those moments yield $\beta(30) = 0.70$ through Equations (9) and (10), which places the crossing of the target reliability index at 17.8 years. The calibrated values are reported in Table 4.

Figure 5 sets out the computational workflow linking environmental generation, degradation evolution, uncertainty propagation, reliability evaluation and decision output.

Table 3. Uncertainty sources, their representation and the pathway through which each affects performance.

Source	Represented by	Primary Pathway
Environmental severity	Ψ, Weibull, Equation (18)	Corrosion rate, stress range, cumulative exposure
Corrosion kinetics	α_c, n_c	Section loss, stress amplification
Corrosion allowance	t_{ca}	Normalisation of corrosion damage
Fatigue loading	$\Delta\sigma_0$, ν	Miner damage accumulation
Fatigue capacity	$\log_{10}\bar{A}$	S–N scatter
Geometry	t_0	Section-loss fraction
Structural capacity	R_0	Limit-state margin
Operational demand	S	Limit-state margin

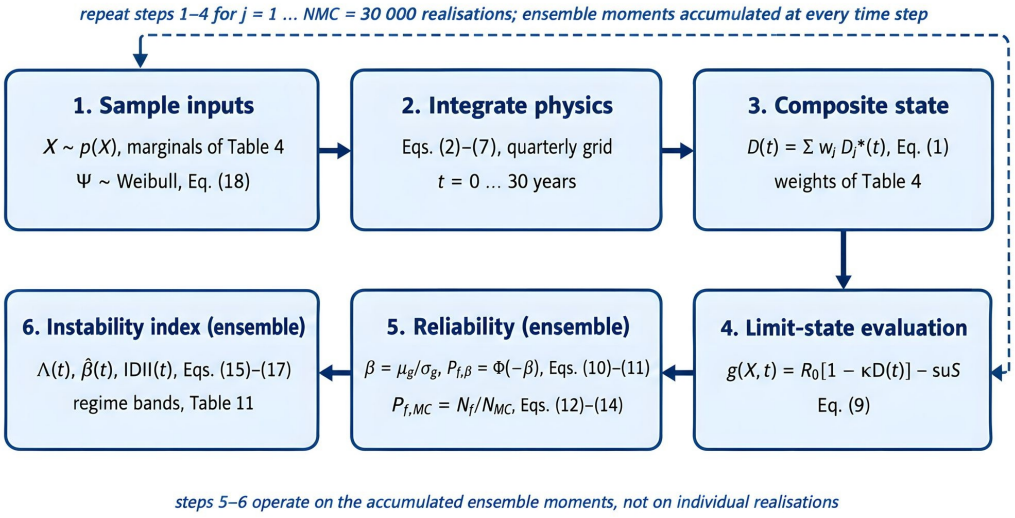


Figure 5. Computational workflow of the probabilistic simulation. Steps 1–4 are performed for each of the 30,000 Monte Carlo realisations, indexed by $r = 1, \ldots, N_{MC}$. The index $j \in \{c, f, e, i\}$ in Step 3 identifies the corrosion, fatigue, environmental-exposure, and interaction components of the composite degradation state, respectively. Ensemble statistics accumulated over the Monte Carlo realisations are subsequently used in Steps 5–6 to evaluate structural reliability and the Infrastructure Degradation Instability Index. The corresponding governing equations are indicated in each workflow block.

Table 4. Complete probabilistic and deterministic parameter specification.

Parameter	Symbol	Distribution	Mean	CoV (Baseline, × σ_E)	Source
Initial wall thickness	t_0	Lognormal	32.0 mm	0.05	Fabrication tolerance
Corrosion allowance	t_{ca}	Lognormal	6.0 mm	0.10	Design specification
Corrosion rate coefficient	α_c	Lognormal	0.120 mm yr^{-n}	0.40 × σ_E	[31,32]
Corrosion time exponent	n_c	Normal, truncated [0.8, 1.6]	1.15	0.08	[31]
Environmental severity	Ψ	Weibull, Equation (18)	1.00	0.563 × σ_E	Table 2
Nominal stress range	$\Delta\sigma_0$	Lognormal	12.5 MPa	0.20 × σ_E	Hot-spot spectral analysis
Cycle rate	ν	Lognormal	4.2 × 10^6 yr^{-1}	0.15	Wave frequency spectrum
S–N intercept	$\log_{10}\bar{A}$	Normal	12.164	s.d. 0.20	Single-slope welded detail
Normalised capacity	R_0	Lognormal	1.00	0.07	Resistance model uncertainty
Normalised demand	S	Lognormal	1.00	0.11 × σ_E	Load model uncertainty
Stress amplification limit	$K_{\sigma,max}$	Deterministic	3.0	—	Concentration factor at maximum admissible section loss
Amplification shape constant	c_K	Deterministic	0.08	—	Section loss at half amplification, Equation (4)
Fatigue–corrosion coupling	a_{fc}	Deterministic	0.35	—	Coupling calibration
Composite weights	w_c, w_f, w_e, w_i	Deterministic	0.38, 0.32, 0.20, 0.10	—	Baseline weighting; sensitivity analysis described in Section 6.7
Capacity transfer coefficient	κ	Deterministic	0.387	—	Calibrated, Section 4.4
Corrosion scaling	c_α	Deterministic	0.54	—	Multiplier in Equation (2); calibrated, Section 4.4
Stress scaling	c_σ	Deterministic	1.21	—	Multiplier in Equation (5); calibrated, Section 4.4
Critical degradation level	D_{crit}	Deterministic	0.70	—	Instability reference
Reference dispersion	$\sigma_{D,ref}$	Deterministic	0.20	—	Accepted end-of-life dispersion
Monte Carlo realisations	N_{MC}	—	30,000	—	Convergence study

Note on the notation σ_E**.** In the CoV column, an entry of the form "0.40 × σ_E" denotes a baseline coefficient of variation of 0.40 at $\sigma_E = 1.0$ that is multiplied by the dimensionless environmental-variability multiplier σ_E when the variability scenario is changed (Section 6.3); for $\sigma_E = 2.0$ the coefficient of variation of α_c is therefore 0.80. σ_E is not a standard deviation and carries no units. Only the four entries containing σ_E (Ψ, α_c, $\Delta\sigma_0$ and S) are scaled; all other entries are unaffected by the multiplier.

5. Data Assimilation and Predictive Analytics

5.1. Sensing Configuration and Data Sources

The assimilation scheme uses two observable quantities, both routinely available on instrumented offshore assets. Wall thickness is measured by ultrasonic testing (UT) at the hot-spot location, with a measurement standard deviation of 0.25 mm consistent with UT deployed from a remotely operated vehicle (ROV) under field conditions. Hot-spot stress range is derived from strain gauges (SG) through rainflow counting, with a standard deviation of 0.80 MPa reflecting gauge accuracy and the spectral processing chain. Both are assimilated annually over the first 20 years of service; the remaining ten years are forecast without further observation, so that the reported performance is genuinely predictive rather than reconstructive.

In an operational deployment the environmental variables would be taken from a co-located wave buoy and anemometer, supplemented by reanalysis fields for periods in which instruments are unavailable; in the present study they are generated from the model of Section 4.3. Two observables were chosen deliberately: the thickness observation constrains the corrosion group $c_\alpha \alpha_c \Psi^{1/2} \langle t - t_i \rangle^{n_c}$, whereas the stress-range observation constrains the fatigue group $K_\sigma c_\sigma \Delta\sigma_0 \Psi^{0.8}$ of Equation (5); only the two together, through their different dependence on Ψ, separate the environmental severity from the mechanism-specific parameters. A single observable leaves one of the two coupled mechanisms unidentified.

5.2. Sequential Bayesian Updating

The posterior over model parameters given observations $\mathbf{z}_{1:k}$ follows Bayes' theorem,

$$p(\mathbf{X} \mid \mathbf{z}_{1:k}) = \frac{p(\mathbf{z}_{1:k} \mid \mathbf{X}) p(\mathbf{X})}{p(\mathbf{z}_{1:k})} \tag{19}$$

Priors. The prior $p(\mathbf{X})$ is the product of the marginals in Table 4. Lognormal priors are adopted for all strictly positive, multiplicatively varying quantities—the corrosion rate coefficient, the stress range, the cycle rate and the thickness variables—because their uncertainty is naturally proportional rather than additive and because a lognormal prior guarantees physical admissibility. A truncated normal is used for the corrosion exponent, which is bounded on physical grounds. A Weibull prior is used for the environmental severity, consistent with the metocean marginals of Table 2. The corrosion rate prior is the widest; its coefficient of variation of 0.40 reflects the scatter reported by Melchers (2003) for immersion conditions and is deliberately wide [32], so that the posterior is data-dominated rather than prior-dominated.

Likelihood. Measurement errors are modelled as independent zero-mean Gaussians, giving

$$p(\mathbf{z}_k \mid \mathbf{X}) \propto exp\left[-\frac{\left(z_k^{UT} - d_c(t_k;\mathbf{X})\right)^2}{2\varsigma_{UT}^2} - \frac{\left(z_k^{SG} - \Delta\sigma_{eff}(t_k;\mathbf{X})\right)^2}{2\varsigma_{SG}^2} \right] \tag{20}$$

with $\varsigma_{UT} = 0.25$ mm and $\varsigma_{SG} = 0.80$ MPa.

Posterior computation. Because Equations (2)–(7) are nonlinear and the priors are non-Gaussian, the posterior is evaluated by sequential importance resampling rather than by a Kalman variant. A population of $N_p = 4000$ particles is propagated through the forward model; at each annual observation the importance weights are updated as

$$\tilde{w}_k^{(j)} \propto w_{k-1}^{(j)}\, p\left(\mathbf{z}_k \mid \mathbf{X}^{(j)}\right), \qquad \sum_j w_k^{(j)} = 1 \tag{21}$$

and the effective sample size is monitored,

$$N_{eff} = \left(\sum_j \left[w_k^{(j)}\right]^2\right)^{-1} \tag{22}$$

when $N_{eff} < N_p/2$ the population is resampled systematically and regularised by multiplicative jitter of coefficient 0.03, which prevents the particle impoverishment that otherwise degrades sequential filters over long observation sequences [43]. Over the 20-year assimilation window the mean effective sample size before resampling was 2019 of 4000 particles, confirming that the population was maintained above the degeneracy threshold at every observation.

Figure 6 illustrates the updating cycle.

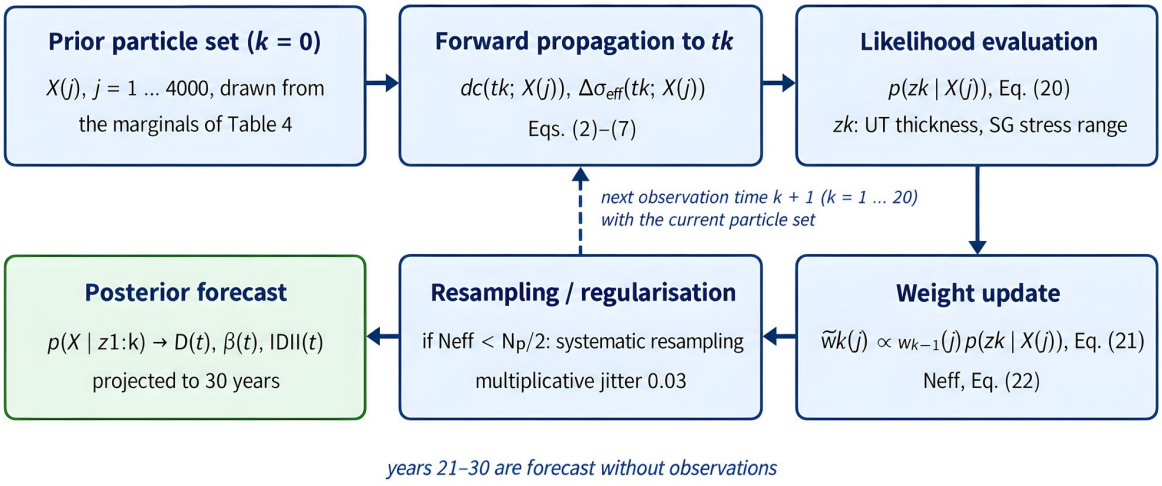


Figure 6. Sequential importance-resampling procedure for Bayesian updating of the coupled degradation model. $\mathbf{X}^{(j)}$ denotes the uncertain parameter vector associated with particle *j*, where $j = 1, \ldots, N_p$ and $N_p = 4000$. The symbol t_k denotes the time of the *k*-th annual observation, with $k = 1, \ldots, 20$. At each observation time, particle-specific predictions are obtained from Equations (2)–(7), and the measurement likelihood is evaluated using Equation (20). Particle weights are updated according to Equation (21), and the effective sample size is determined using Equation (22). Systematic resampling and regularisation are applied when the effective sample size falls below $N_p/2$. Following assimilation of the final annual observation, the posterior particle ensemble is propagated to the 30-year assessment horizon without further measurement updates.

5.3. Physics-Constrained Surrogate Model

The surrogate accelerates repeated evaluation of the degradation model in scenario and design studies that query the model many times within its training envelope; the particle filter of Section 5.2 and all lifecycle forecasts reported in this paper use the closed-form model of Section 3.2 directly. Its role is bounded: it interpolates within the trained envelope, and it is not used to extrapolate in time. Section 6.8 gives the evidence for that restriction.

Following the missing-physics decomposition of Dourado and Viana (2020), the surrogate does not predict the degradation state directly [49]. It predicts only the part of the state that the analytical model does not already supply:

$$\widehat{D}(t,\mathbf{X}) = \underbrace{w_c \min\left(\frac{c_\alpha \alpha_c \Psi^{1/2} \langle t - t_i \rangle^{n_c}}{t_{ca}}, 1\right) + w_e \min\left(\frac{t}{T_d}\Psi, 1\right)}_{\text{known physics, } D_{phys}} + softplus\left[\mathcal{N}_{\boldsymbol{\Phi}}(t,\mathbf{X})\right] \tag{23}$$

The known term carries the corrosion and environmental-exposure contributions, which are directly measurable. The network learns the fatigue and interaction contributions, which are not. Two physical properties are thereby imposed by construction rather than by penalty: the softplus output is non-negative, so the network can only add damage, and $D_{phys}(0) = 0$ with a non-negative residual guarantees $\widehat{D}(0) = 0$. Monotonicity in time and the upper bound $\widehat{D} \leq 1$ are not imposed by the architecture; both were verified numerically on the test set rather than assumed.

The surrogate model specifications are summarised in Table 5.

Table 5. Surrogate model specification.

Item	Specification
Inputs	α_c, n_c, t_{ca}, $\Delta\sigma_0$, Ψ, t (6 inputs, min–max normalised to [−1,1]); ν, $\log_{10}\bar{A}$ and t_0 held at their Table 4 means
Output	Residual damage, added to D_{phys} per Equation (23)
Hidden layers	2 × 48 units
Activation	tanh (hidden), softplus (output)
Trainable parameters	2737
Loss	Mean squared error on the residual target
Optimiser	Adam, learning rate 4×10^{-3}
Epochs / batch size	400 / 128
Training set	3000 degradation trajectories (121 time points each), sampled from the prior
Test set	1200 independent trajectories
Inference cost	0.68 µs per evaluation

Training and test sets are drawn from the same prior; the test trajectories are never seen during training. Accuracy is reported through root mean square error, mean absolute error, mean absolute percentage error and the coefficient of determination:

$$RMSE = \sqrt{\frac{1}{n}\sum_i \left(\hat{D}_i - D_i\right)^2}, \quad MAE = \frac{1}{n}\sum_i \left|\hat{D}_i - D_i\right| \tag{24}$$

$$MAPE = \frac{100}{n}\sum_i \left|\frac{\hat{D}_i - D_i}{D_i}\right|, \quad R^2 = 1 - \frac{\sum_i \left(\hat{D}_i - D_i\right)^2}{\sum_i (D_i - \bar{D})^2} \tag{25}$$

MAPE is evaluated only for $D_i > 0.02$ to avoid division by near-zero early-life values. Results are given in Section 6.8, benchmarked against an unconstrained network of identical size trained on the same data.

5.4. Global Sensitivity Analysis

Variance-based sensitivity indices are computed by the Saltelli estimator with a base sample of 4096, giving 49,152 model evaluations for ten inputs. First-order and total-order indices are

$$S_i = \frac{Var_{X_i}[E(Y \mid X_i)]}{Var(Y)}, \qquad S_{Ti} = 1 - \frac{Var_{\mathbf{X}_{\sim i}}[E(Y \mid \mathbf{X}_{\sim i})]}{Var(Y)} \tag{26}$$

evaluated for two outputs: the end-of-life degradation state $D(30)$ and the limit-state margin $g(30)$.

5.5. Computational Cost and Real-Time Capability

The framework is not a real-time controller and does not need to be. The physical processes it tracks evolve over months to years, and the observations that drive it arrive annually to quarterly. The relevant performance requirement is therefore that a complete state update should be fast relative to the monitoring interval, which it comfortably is: a full sequential assimilation over 20 annual observation times with 4000 particles over a 30-year horizon completes in 0.62 s per asset on a single CPU core. Surrogate inference costs 0.68 µs per evaluation. A fleet of 1000 instrumented assets could therefore be updated in approximately ten minutes of single-core time, or seconds on a modest cluster.

The genuinely time-critical elements of an operational deployment—data acquisition, signal conditioning, rainflow counting and anomaly flagging—run at the edge on the sensing hardware and are not part of the inference loop described here. This separation between fast edge processing and slower physics-based state estimation is what makes the architecture practical; Section 7.4 discusses how it scales when the forward model is a finite-element system rather than the reduced formulation used here.

6. Results and Discussion

6.1. Degradation Evolution

Figure 7 shows the degradation state over the 30-year horizon. The mean rises to 0.636 at end of life, with a standard deviation of 0.300 and a 90% credible interval spanning 0.126 to 0.990. The dispersion grows more slowly than the mean in absolute terms and much more slowly in relative terms: the coefficient of variation falls from 1.02 at five years to 0.47 at thirty. Early-life relative scatter is large because the scatter of the environmental severity and of the stress range, which enter the fatigue rate through $\Psi^{2.4}$ and $\Delta\sigma_0^3$, acts on a mean that is still small; as damage accumulates the outcome becomes more determinate in relative terms while the absolute spread continues to widen.

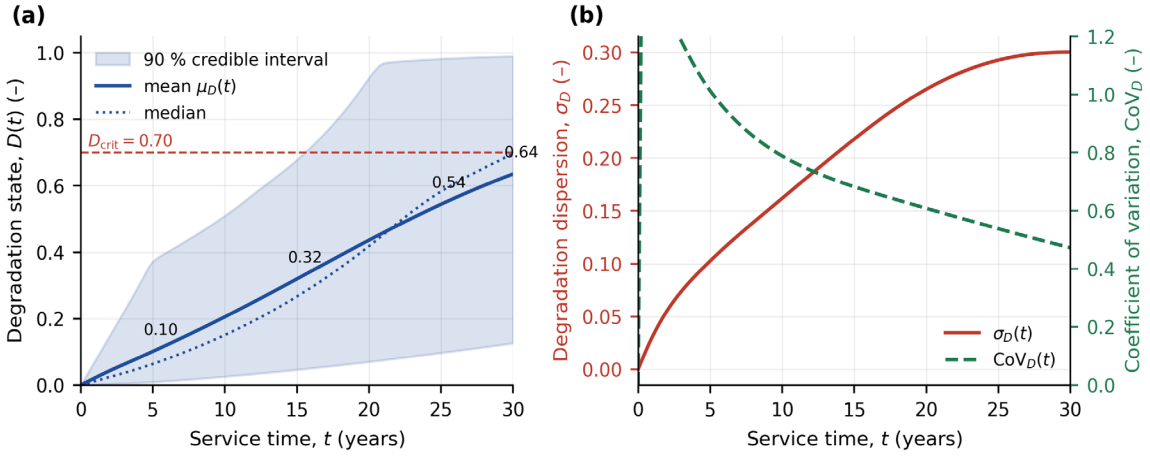


Figure 7. (**a**) Degradation state with mean, median and 90% credible interval, and the critical level D_{crit} = 0.70. (**b**) Absolute dispersion and coefficient of variation.

The degradation rate is not monotonic. It falls slightly from 0.0221 yr^{-1} at commissioning, when only exposure and fatigue accumulate, to 0.0199 yr^{-1} at five years, rises to a maximum of 0.0234 yr^{-1} at fifteen years once corrosion-driven stress amplification is established, then falls to 0.0165 yr^{-1} at thirty. The peak marks the point at which corrosion-driven stress amplification is accelerating fatigue most strongly; beyond it, the saturation of the individual damage components in Equation (1)—the corrosion allowance is exhausted, Miner damage approaches unity, cumulative exposure approaches its bound—slows the composite index even though the underlying physical deterioration continues. This is an important qualification on the common statement that offshore degradation "accelerates in late life": in a bounded composite index the acceleration appears in the *reliability margin* and in the *dispersion*, not necessarily in the mean index itself. The behaviour of $\beta(t)$ in Section 6.2 makes this explicit.

Component contributions at 30 years are 31.8% corrosion, 35.7% fatigue, 24.2% environmental exposure and 8.3% interaction (Figure 2c). The interaction component carries the smallest weight but records the coupling state of the two mechanisms: it tracks the stress amplification through which corrosion feeds into fatigue (Equation (5)), while the fatigue-to-corrosion feedback acts through a_{fc} in Equation (2). A supplementary run with the fatigue-to-corrosion coupling set to zero and the interaction weight redistributed reduced $D(30)$ from 0.636 to 0.549 and delayed the crossing of β_{target} from 17.8 to 20.5 years. Neglecting the feedback therefore overstates qualified service life by 2.7 years, or 15%.

6.2. Reliability Evolution

The reliability index falls from 3.48 at commissioning to 0.70 at 30 years, crossing the target value of 1.5 at 17.8 years (Figure 8). Failure probability rises correspondingly from 2.51×10^{-4} to 2.42×10^{-1}, and reliability falls from 0.9997 to 0.7580. Table 6 gives the lifecycle quantities at five-year intervals; every (β, P_f, R) triple satisfies $P_{f,\beta} = \Phi(-\beta)$ and $R = 1 - P_{f,\beta}$.

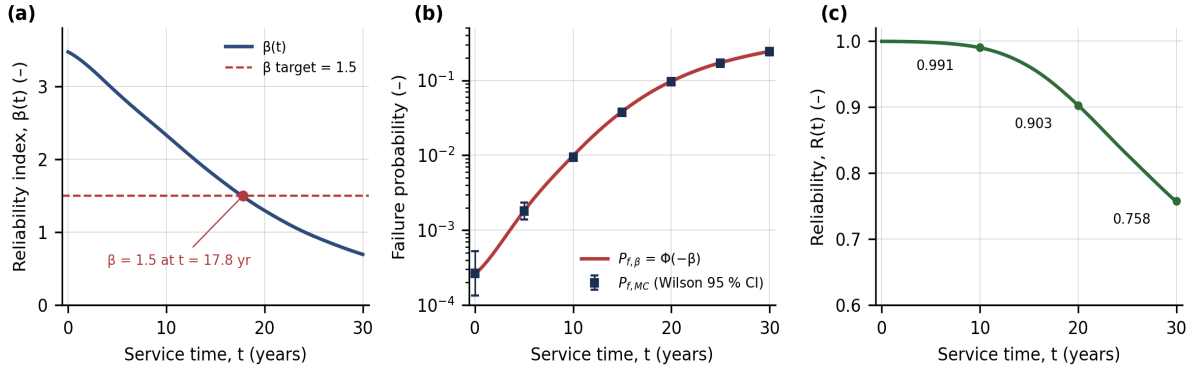


Figure 8. (**a**) Reliability index with the target level and its crossing time. (**b**) Two-moment failure probability $P_{f,\beta}(t)$ (line) with the empirical Monte Carlo estimates $P_{f,MC}$ of Table 7 and their Wilson 95% intervals (markers). (**c**) Reliability function.

Table 6. Quantitative lifecycle assessment results for the reference component (σ_E = 1.0).

t (yr)	μ_D	σ_D	CoV(D)	5th–95th pct. of D	$\dot{D}$ (yr^{-1})	β	$P_{f,\beta}$	R	Λ	*IDII*	Regime
0	0.000	0.000	—	0.000–0.000	0.0221	3.48	2.51×10^{-4}	0.9997	0.000	0.000	Stable
5	0.101	0.103	1.02	0.009–0.372	0.0199	2.92	1.75×10^{-3}	0.9982	0.437	0.183	Controlled
10	0.206	0.162	0.79	0.026–0.510	0.0219	2.35	9.39×10^{-3}	0.9906	0.761	0.327	Controlled
15	0.320	0.218	0.68	0.046–0.678	0.0234	1.78	3.75×10^{-2}	0.9625	1.096	0.480	Transitional
20	0.436	0.265	0.61	0.070–0.925	0.0228	1.30	9.68×10^{-2}	0.9032	1.298	0.599	Unstable
25	0.544	0.293	0.54	0.097–0.982	0.0201	0.95	1.71×10^{-1}	0.8289	1.261	0.666	Critical
30	0.636	0.300	0.47	0.126–0.990	0.0165	0.70	2.42×10^{-1}	0.7580	1.059	0.695	Critical

Stage-averaged values are β = 2.92, 1.80 and 0.97 for the three ten-year stages, with corresponding mean degradation states of 0.102, 0.320 and 0.541 and mean *IDII* of 0.180, 0.474 and 0.659. The reliability index falls almost linearly through the first two stages and then flattens as the composite index saturates, whereas the failure probability rises throughout, with its largest five-year increments in the final two stages.

Empirical verification of the failure probability. Because the complete set of limit-state samples is available, the empirical failure probability of Equation (12) was counted directly at each evaluation year and compared with the two-moment value, as reported in Table 7. The absolute difference of Equation (13) is at most 2.2×10^{-4}, and the two-moment value lies inside the Wilson 95% interval of the empirical estimate at every year. At commissioning the empirical estimate rests on only 8 failed realisations out of 30,000, so its interval (1.35×10^{-4} to 5.26×10^{-4}) is wide and the relative difference is 6%; from year 5 onward the relative difference is below 3%, and from year 10 onward below 1%. The agreement holds despite the strong skewness of $D(t)$ evident in Table 6 (coefficient of variation up to 1.02, asymmetric credible intervals): the limit state of Equation (9) is linear in D, and roughly half of the variance of the margin g arises from the near-lognormal capacity and demand terms (as quantified by the variance decomposition in Section 6.4), so the distribution of g is far less skewed than that of D in the region that governs the failure probability. The adequacy of the two-moment approximation for this limit state is therefore established empirically rather than assumed; $P_{f,\beta}$ is used for the remaining results, and the empirical values are retained in Table 7 as the reference against which that adequacy has been demonstrated.

Table 7. Empirical Monte Carlo failure probability against the two-moment value, N_{MC} = 30,000 realisations, σ_E = 1.0. CI = Wilson 95% confidence interval of $P_{f,MC}$, Equation (14); ΔP_f is the absolute difference of Equation (13) between $P_{f,MC}$ and the exact value of $\Phi(-\beta)$ for the tabulated β; $\delta P_f = \Delta P_f / P_{f,MC}$.

t (yr)	β	$P_{f,\beta} = \Phi(-\beta)$	N_f / N_{MC}	$P_{f,MC}$	95% CI of $P_{f,MC}$	ΔP_f	δP_f (%)
0	3.48	2.51×10^{-4}	8 / 30,000	2.67×10^{-4}	1.35×10^{-4}–5.26×10^{-4}	1.6×10^{-5}	6.0
5	2.92	1.75×10^{-3}	54 / 30,000	1.80×10^{-3}	1.38×10^{-3}–2.35×10^{-3}	5.0×10^{-5}	2.8
10	2.35	9.39×10^{-3}	284 / 30,000	9.47×10^{-3}	8.43×10^{-3}–1.06×10^{-2}	8.0×10^{-5}	0.84
15	1.78	3.75×10^{-2}	1128 / 30,000	3.76×10^{-2}	3.55×10^{-2}–3.98×10^{-2}	6.2×10^{-5}	0.16
20	1.30	9.68×10^{-2}	2900 / 30,000	9.67×10^{-2}	9.34×10^{-2}–1.00×10^{-1}	1.3×10^{-4}	0.14
25	0.95	1.71×10^{-1}	5125 / 30,000	1.71×10^{-1}	1.67×10^{-1}–1.75×10^{-1}	2.2×10^{-4}	0.13
30	0.70	2.42×10^{-1}	7260 / 30,000	2.42×10^{-1}	2.37×10^{-1}–2.47×10^{-1}	3.6×10^{-5}	0.02

6.3. Influence of Environmental Variability

Table 8 and Figure 9 report the effect of the environmental variability multiplier at fixed mean severity, with σ_E scaling the environment-driven coefficients of variation of Ψ, α_c, $\Delta\sigma_0$ and S together, as defined in Section 4.3. The result is the principal finding of the study.

Table 8. Effect of environmental variability at constant mean environmental severity.

σ_E	$\mu_D(30)$	$\sigma_D(30)$	$\beta(20)$	$P_{f,\beta}(20)$	$\beta(30)$	$IDII(30)$	Year at Which β = 1.5
0.5	0.728	0.209	1.68	4.65×10^{-2}	0.64	0.700	21.5
1.0	0.636	0.300	1.30	9.68×10^{-2}	0.70	0.695	17.8
1.5	0.549	0.341	1.15	1.25×10^{-1}	0.76	0.657	14.5
2.0	0.478	0.361	1.06	1.45×10^{-1}	0.78	0.617	10.8

Quadrupling the variability multiplier reduces the *mean* end-of-life degradation state by 34%, from 0.728 to 0.478, while increasing its standard deviation by 73%, from 0.209 to 0.361. The reliability index at twenty years falls by 37%, from 1.68 to 1.06, and the year in which the target reliability is first breached advances from 21.5 to 10.8—a loss of more than ten years of qualified service life.

The mean degradation state falls because the environmental severity index is normalised to unit mean; increasing σ_E at fixed mean redistributes probability mass toward low severity with a compensating heavy upper tail, and the bounded damage components saturate under the tail rather than averaging linearly through it. Reliability nevertheless deteriorates through the first two decades, because β depends on the ratio of the mean margin to its standard deviation and the variance contributed by the tail grows faster than the mean margin recovers; $\beta(20)$ falls from 1.68 to 1.06 and the target is breached a decade earlier. Only very late in life do the curves cross (Table 8: $\beta(30)$ rises from 0.64 to 0.78 with σ_E, because the lower mean degradation then outweighs the added variance), by which time the target has long been breached in every case; the late-life crossover therefore has no bearing on the qualified service life.

The practical consequence is direct. Assessing climate change effects on offshore assets by asking how much the *mean* wave climate will shift is insufficient, and can be actively misleading: an assessment based on mean degradation would conclude that the high-variability case is the more benign one, when in fact it loses a decade of service life. Variability must be propagated as a distribution, not summarised as a mean.

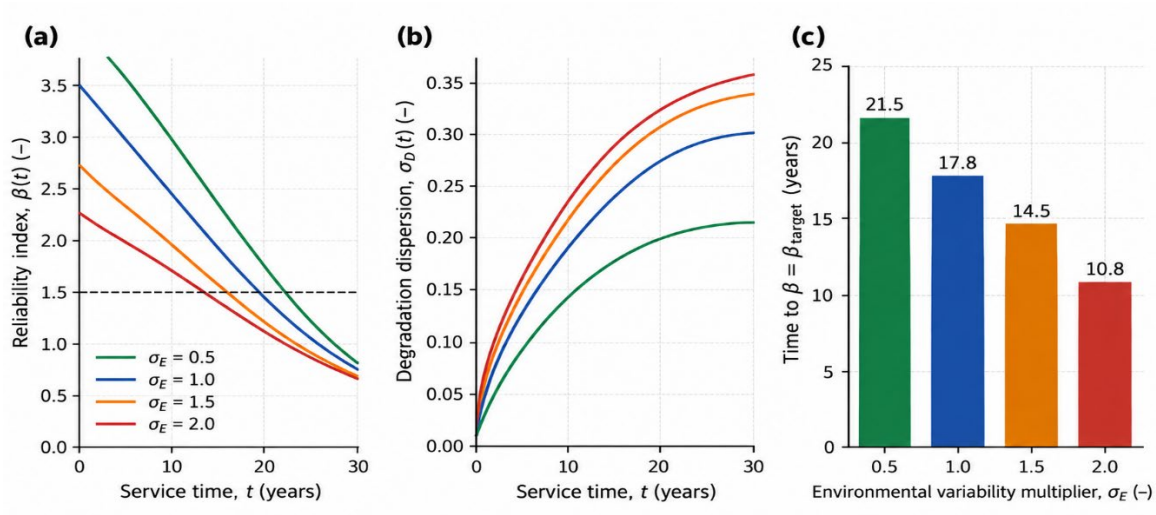


Figure 9. Effect of environmental variability at constant mean severity. (**a**) Evolution of the structural reliability index for different levels of the environmental variability multiplier, with the target reliability index indicated by the horizontal reference line. (**b**) Evolution of the dispersion of the composite degradation state. (**c**) Year at which the target reliability index of 1.5 is first reached as a function of the dimensionless environmental variability multiplier, σ_E.

6.4. Global Sensitivity

Table 9 gives the Sobol' indices. For the end-of-life degradation state, environmental severity accounts for 86.8% of the total-order variance, the corrosion rate coefficient for 11.2% and the nominal stress range for 4.4%; all remaining inputs contribute below 3% each. The near-equality of first-order and total-order indices indicates that interaction effects between

inputs are modest at this level, so the variance is essentially additive in the dominant parameters (Figure 10). Where a first-order index exceeds its total-order counterpart (Ψ and t_{ca} for $D(30)$, S for $g(30)$), and where the first-order indices for $D(30)$ sum to slightly more than unity (1.024), the discrepancies are within the Monte Carlo sampling error of the Saltelli estimator at the base sample used, and the affected pairs should be read as equal.

Table 9. First-order (S_i) and total-order (S_{Ti}) Sobol' indices.

Input	S_i for $D(30)$	S_{Ti} for $D(30)$	S_i for $g(30)$	S_{Ti} for $g(30)$
Environmental severity, Ψ	0.893	0.868	0.408	0.454
Corrosion rate, α_c	0.083	0.112	0.049	0.058
Stress range, $\Delta\sigma_0$	0.034	0.044	0.010	0.023
S–N intercept, $\log_{10}\bar{A}$	0.007	0.028	0.008	0.014
Corrosion allowance, t_{ca}	0.007	0.004	0.002	0.002
Operational demand, S	0.000	0.000	0.298	0.295
Initial capacity, R_0	0.000	0.000	0.178	0.179

For the limit-state margin the picture changes, as it must: environmental severity remains the largest single contributor at 45.4%, but demand and capacity model uncertainty together account for 47.4%. Degradation modelling and load–resistance modelling are therefore of comparable importance to reliability, and refining one without the other yields limited benefit. This provides a defensible basis for allocating characterisation effort: for predicting *damage*, metocean characterisation dominates; for predicting *reliability*, load and resistance model uncertainty matter equally.

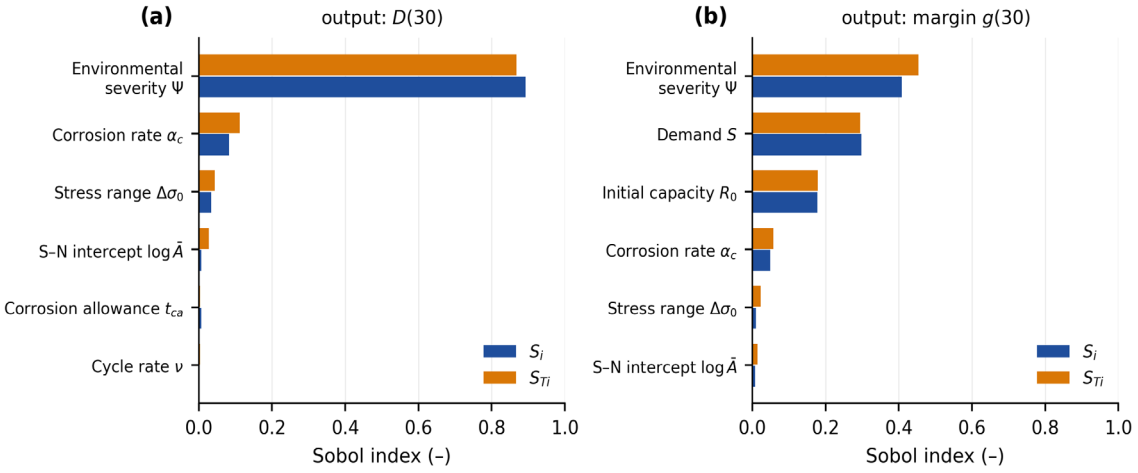


Figure 10. Global variance decomposition. (**a**) Indices for the end-of-life degradation state. (**b**) Indices for the limit-state margin.

6.5. Performance of Bayesian Updating

Table 10 reports the assimilation benchmark across 150 independent synthetic assets, each with its own true parameter set and its own realised measurement noise. Reporting a distribution over assets rather than a single illustrative trajectory is deliberate: the improvement obtained on one favourable realisation is not evidence of method performance.

Table 10. Sequential Bayesian updating versus non-updating prediction, over 150 synthetic assets.

Metric	Value
Reduction in predictive dispersion $\sigma_D(30)$	81.5% (IQR 78.3–85.1%)
Reduction in 90% credible-interval width at 30 yr	80.6%
Reduction in absolute prediction error at 30 yr (median)	80.9%
Mean effective sample size	2019 of 4000 particles
Assimilation cost per asset	0.62 s (single CPU core)
Observations assimilated	20 annual thickness + 20 annual stress-range (20 observation times)

The filter reduces the predictive standard deviation of the 30-year degradation state by 81.5%, with an interquartile range of 78.3% to 85.1% across assets. Because the final ten years are forecast without observation, this is a genuine forecasting gain rather than a smoothing artefact. Figure 11 shows a representative asset: the prior predictive interval widens continuously, whereas the posterior tracks the reference trajectory and its interval remains narrow after the observation window closes at year 20.

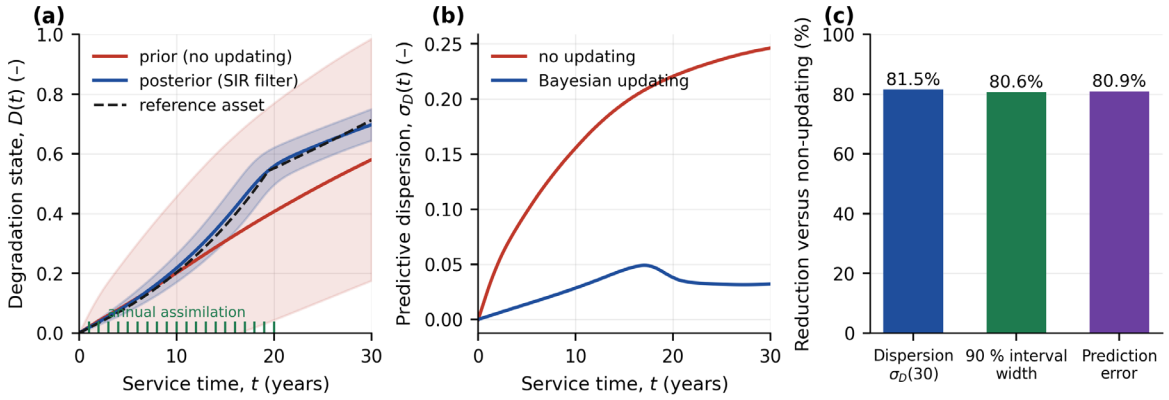


Figure 11. Sequential Bayesian updating for a representative asset. (**a**) Prior and posterior degradation predictions with 90% intervals against the reference trajectory; assimilation ends at year 20. (**b**) Predictive dispersion. (**c**) Reduction achieved across 150 assets.

The mean effective sample size of 2019 of 4000 particles confirms that the resampling and regularisation strategy of Section 5.2 prevented degeneracy. An earlier implementation using batch importance sampling over all 20 observations simultaneously collapsed to an effective sample size of 20, and produced posterior dispersions that were not trustworthy; the comparison is reported because this failure mode is not visible when the posterior computation method is left unstated.

6.6. Instability Index Behaviour and Threshold Calibration

IDII rises monotonically from zero to 0.695 over the 30-year design life in the baseline scenario (Figure 3). The influence of environmental variability on the evolution of IDII is examined separately in Figure 12. Its two constituents behave differently: the instability driver Λ rises from zero, peaks at 1.30 near twenty years and then declines slightly as the degradation rate falls, while the normalised reliability margin $\hat{\beta}$ falls steadily from 2.32 to 0.47. Their crossover at 15.7 years corresponds to *IDII* = 0.5 and marks the point at which the accumulation of damage and uncertainty first outweighs the margin remaining.

Thresholds are calibrated by asking which reliability index corresponds to a given *IDII* level, rather than by assertion. Table 11 gives the result.

Table 11. IDII regime bands, their reliability correspondence for the reference component, and the associated management action.

Regime	*IDII* Range	β at Upper Bound	Year Reached	Management Action
Stable	<0.15	3.06	3.8	Baseline inspection interval
Controlled	0.15–0.35	2.26	10.7	Routine monitoring, no change
Transitional	0.35–0.50	1.71	15.7	Shorten inspection interval; review corrosion protection
Unstable	0.50–0.65	1.05	23.4	Detailed inspection; plan intervention
Critical	≥0.65	—	—	Intervention or reassessment before continued operation

The correspondence is useful in both directions. An operator who measures *IDII* = 0.5 can infer that the reliability index has fallen to approximately 1.7 and that the target level of 1.5 will be reached within roughly two years at the current rate. Conversely, the reliability target itself maps to an IDII trigger—for this component, $\beta = 1.5$ corresponds to *IDII* ≈ 0.54–0.55—which can be monitored directly without recomputing the full limit state at every inspection.

Two cautions apply. The numerical correspondence in Table 11 is specific to the reference component, which was configured deliberately as an aggressive demonstration case (Section 4.2), and to its target reliability; the *procedure* transfers, the *numbers* must be recalibrated per asset class, and the crossing years are not representative of code-compliant assets in normal service. And IDII is sensitive to environmental variability: at $\sigma_E = 2.0$ the index reaches only 0.617 at thirty years compared with 0.695 at $\sigma_E = 1.0$ (Table 8), because the mean degradation rate is lower even though the target reliability is breached seven years earlier. IDII should therefore be read alongside β, not instead of it; the seven-year lead time between the transitional threshold and the target quoted in Section 7.2 applies to the baseline scenario, and at higher variability both the threshold and the target are reached earlier (Figure 12a, Table 8) and the lead shortens. Figure 12b shows the mapping between degradation state and IDII coloured by reliability index, which makes the joint reading explicit.

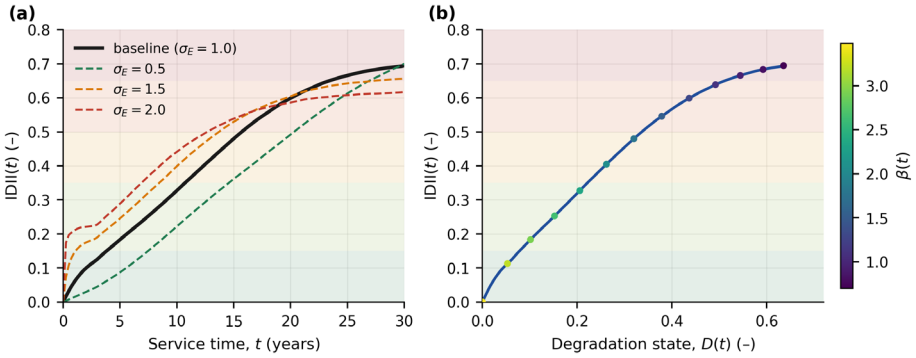


Figure 12. (**a**) IDII evolution for four environmental variability levels, with regime bands. (**b**) Relationship between degradation state and IDII, coloured by the corresponding reliability index, with the regime bands of (**a**).

6.7. Sensitivity of IDII to the Composite-Index Weights

The weights of Equation (1) are assigned on engineering grounds rather than derived from data, and a reader is entitled to ask how much of the decision output depends on that assignment. Because $D(t)$ is linear in the weights, every weight set can be evaluated exactly from the component statistics of the baseline simulation of Section 4.4: the means of the four components and their covariance matrix at every point of the quarterly grid determine $\mu_D(t)$ and $\sigma_D(t)$ for any weight vector, from which $\beta(t)$ follows through Equations (9) and (10) and *IDII* through Equations (15)–(17), with all other parameters held at the values of Table 4. Twenty-one weight sets were evaluated in this way: the baseline; sixteen one-at-a-time (OAT) cases in which each weight in turn is reduced or increased by 20% and by 30% while the remaining three are renormalised so that the weights still sum to unity; and four stress-test weightings that deliberately favour one mechanism (equal weighting; corrosion-, fatigue- and environment-dominant). For each case the end-of-life moments of D, *IDII*(30), the *IDII* threshold crossing years and the year at which β = 1.5 is reached were recorded. The results are given in Table 12 and Figure 13.

Table 12. Sensitivity of the end-of-life degradation moments, *IDII*(30), the *IDII* threshold-crossing years and the reliability-target year to the composite-index weights, σ_E = 1.0. OAT = one weight perturbed by the stated percentage with the other three renormalised to unit sum (weights rounded to three decimals); the four stress-test rows are engineering weightings that favour one mechanism; because of the rounding, some rows sum to 0.999 or 1.001. $t_{0.35}$, $t_{0.50}$ and $t_{0.65}$ are the years at which *IDII* first reaches the transitional, unstable and critical thresholds; $t_{\beta=1.5}$ is the year at which the target reliability index is reached. ">30" denotes a threshold not reached within the 30-year horizon.

Case	w_c	w_f	w_e	w_i	$\mu_D(30)$	$\sigma_D(30)$	*IDII*(30)	$t_{0.35}$ (yr)	$t_{0.50}$ (yr)	$t_{0.65}$ (yr)	t_β = 1.5 (yr)
Baseline	0.38	0.32	0.20	0.10	0.636	0.300	0.695	10.7	15.7	23.4	17.8
Corrosion − 30%	0.266	0.379	0.237	0.118	0.655	0.307	0.712	9.0	14.0	21.4	16.5
Corrosion − 20%	0.304	0.359	0.225	0.112	0.649	0.304	0.706	9.6	14.6	22.1	17.0
Corrosion + 20%	0.456	0.281	0.175	0.088	0.623	0.297	0.686	11.7	16.7	24.7	18.6
Corrosion +30%	0.494	0.261	0.163	0.082	0.617	0.295	0.682	12.3	17.2	25.4	19.0
Fatigue − 30%	0.434	0.224	0.228	0.114	0.626	0.285	0.670	12.1	17.2	26.3	18.7
Fatigue − 20%	0.416	0.256	0.219	0.109	0.629	0.290	0.678	11.6	16.7	25.3	18.4
Fatigue + 20%	0.344	0.384	0.181	0.091	0.643	0.311	0.712	9.7	14.7	21.9	17.2
Fatigue + 30%	0.326	0.416	0.172	0.086	0.646	0.316	0.721	9.2	14.1	21.2	16.8
Environment − 30%	0.409	0.344	0.140	0.108	0.626	0.303	0.693	10.9	16.0	23.7	18.1
Environment − 20%	0.399	0.336	0.160	0.105	0.629	0.302	0.694	10.9	15.9	23.6	18.0
Environment + 20%	0.361	0.304	0.240	0.095	0.643	0.298	0.696	10.5	15.5	23.1	17.6
Environment + 30%	0.352	0.296	0.260	0.092	0.646	0.298	0.697	10.5	15.4	23.0	17.4
Interaction − 30%	0.393	0.331	0.207	0.070	0.640	0.305	0.706	10.5	15.4	22.7	17.6
Interaction − 20%	0.388	0.327	0.204	0.080	0.638	0.303	0.702	10.6	15.5	22.9	17.7
Interaction + 20%	0.372	0.313	0.196	0.120	0.634	0.297	0.688	10.8	15.9	24.0	17.9
Interaction + 30%	0.367	0.309	0.193	0.130	0.632	0.295	0.684	10.9	16.0	24.3	18.0
Equal weighting	0.25	0.25	0.25	0.25	0.635	0.272	0.643	11.0	16.6	>30	18.0
Corrosion dominant	0.50	0.25	0.15	0.10	0.612	0.292	0.673	12.6	17.6	26.2	19.3
Fatigue dominant	0.25	0.50	0.15	0.10	0.656	0.328	0.738	7.8	12.7	19.5	15.9
Environment dominant	0.25	0.25	0.40	0.10	0.671	0.292	0.703	9.6	14.3	21.8	16.4

Three observations follow. First, the direction of every shift is fixed by the component statistics. The fatigue component has the largest end-of-life dispersion of the four (standard deviation 0.435, against 0.305, 0.303 and 0.164 for the corrosion, environmental-exposure and interaction components) and a mean (0.708) above that of the composite, so increasing w_f raises both μ_D and σ_D, lowers β and raises *IDII*, and brings every threshold forward; the corrosion component has a mean (0.531) below that of the composite and a smaller dispersion (0.305), so increasing w_c has the opposite effect and delays every threshold. The environmental-exposure component has the highest mean (0.768) and a dispersion close to that of corrosion, so increasing w_e raises μ_D but, because the renormalisation reduces the high-dispersion fatigue weight, slightly lowers σ_D (0.298 against 0.300 at +30%); the two effects on β partly cancel and the net shift is small. The interaction component has the lowest mean (0.527) and by far the smallest dispersion (0.164), and its weight is the smallest of the four, so it has the least influence. In the OAT cases the year at which β = 1.5 is reached moves by at most 0.8 years for ±20% and 1.3 years for ±30% (both extremes from the corrosion weight, followed by the fatigue weight), and the transitional threshold year $t_{0.35}$ by at most 1.1 and 1.7 years; at ±30% the environmental weight moves every crossing year by 0.4 years or less, and the interaction weight moves the transitional and reliability-target years by 0.2 years or less, the unstable threshold year by 0.3 years or less and the critical threshold year by at most 0.9 years.

Second, the stress tests bound the decision output. Across all 21 cases the year at which β = 1.5 is reached lies between 15.9 years (fatigue-dominant) and 19.3 years (corrosion-dominant), $t_{0.35}$ between 7.8 and 12.6 years, and *IDII*(30) between 0.643 and 0.738. The ordering of the regime transitions is the same in every case and the horizon ends in the critical band in 20 of the 21 cases; the exception is equal weighting, which assigns the largest relative weight to the low-dispersion interaction component, reduces $\sigma_D(30)$ from 0.300 to 0.272 and ends the horizon at *IDII*(30) = 0.643, in the upper part of the unstable band, with the critical threshold not reached within 30 years.

Third, and most important for maintenance decisions, the warning lead time that the framework provides is stable. The interval between the transitional threshold (*IDII* = 0.35, the point at which the inspection interval is shortened) and the year at which β = 1.5 is reached is 7.1 years for the baseline and lies between 6.6 years (fatigue −30%) and 8.1 years (fatigue-dominant) across the 21 cases; it exceeds 6.5 years in every case. The maintenance policy of Section 7.2, which keys inspection decisions to the IDII regime bands, therefore issues the same sequence of actions with a lead time that varies by no more than ±1 year around the baseline value under every weighting examined. The weights govern the absolute level and timing of the composite index; they do not alter the decision logic it supports. The

thresholds of Table 11 nevertheless remain a calibration for the stated baseline weights, and the recalibration procedure of Section 6.6 should be repeated if a different weighting is adopted for another asset class.

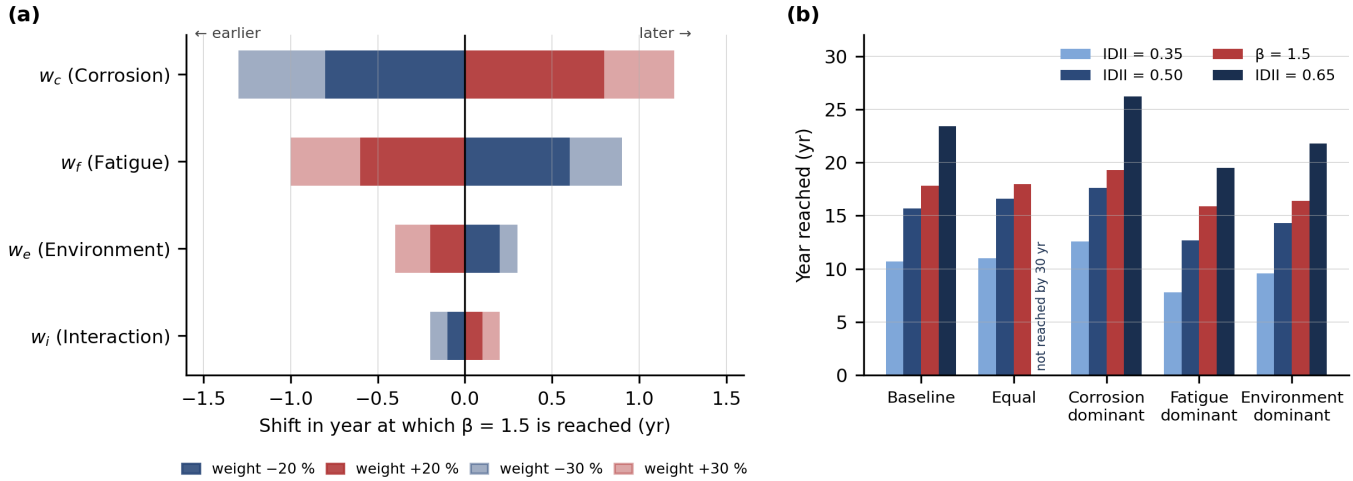


Figure 13. Sensitivity of the reliability-target year and the IDII threshold years to the composite-index weights. (**a**) Shift in the year at which $\beta = 1.5$ is reached for one-at-a-time perturbations of ±20% and ±30% in each weight, with the remaining weights renormalised. (**b**) Years at which *IDII* reaches 0.35, 0.50 and 0.65 and at which $\beta = 1.5$ is reached, for the baseline and the four stress-test weightings of Table 12; for equal weighting the critical threshold is not reached within the 30-year horizon.

6.8. Surrogate Accuracy and the Limits of Extrapolation

Table 13 reports surrogate accuracy against training set size, benchmarked against an unconstrained network of identical size and training budget. Results are averaged over three random initialisations (Figure 14).

Table 13. Surrogate accuracy on 1200 held-out trajectories. PC = physics-constrained residual model, Equation (23); UC = unconstrained network.

N_{tr}	PC RMSE	PC MAE	PC MAPE (%)	PC R^2	UC RMSE	UC MAE	UC MAPE (%)	UC R^2
40	0.0926	0.0744	47.22	0.865	0.0705	0.0501	26.11	0.922
80	0.0833	0.0673	43.78	0.891	0.0605	0.0443	25.82	0.942
160	0.0678	0.0537	36.36	0.926	0.0441	0.0318	17.86	0.969
400	0.0464	0.0320	18.84	0.966	0.0383	0.0272	17.36	0.977
1000	0.0131	0.0088	4.89	0.997	0.0133	0.0089	4.52	0.997
3000	0.0078	0.0051	2.87	0.999	0.0088	0.0058	3.13	0.999

At the full training set the physics-constrained model achieves RMSE 0.0078, MAE 0.0051, MAPE 2.87% and $R^2 = 0.999$, marginally better than the unconstrained network on every metric. At small training sets the unconstrained network is more accurate, because the residual formulation of Equation (23) inherits the approximation error of the analytical term and has less freedom to compensate for it. Both models satisfied monotonicity in time and the upper bound of unity at every tested point; only the constrained model, however, guarantees non-negativity and a zero initial state by construction, which matters when the surrogate is queried inside an optimisation loop that can seek out pathological regions.

The physics constraint delivers structural guarantees and a modest accuracy benefit in the data-rich regime; it does not deliver a large accuracy advantage, and claiming otherwise would not be supported by the evidence.

The extrapolation test is more consequential. Training on trajectories truncated at 12 years and testing on the interval beyond 20 years, both models fail: the coefficient of determination is negative in five of six configurations for the constrained model and four of six for the unconstrained one, meaning both are worse than predicting the test-set mean. The single favourable case ($N_{tr} = 1000$, $R^2 = 0.910$ constrained against 0.756 unconstrained) is not reproducible across training set sizes and should not be treated as evidence of extrapolation capability.

This result determines the architecture. Neural surrogates—physics-constrained or otherwise—cannot be relied upon to forecast degradation into a time regime they have not seen, and the constraints imposed in Equation (23) are not sufficient to change that. The framework therefore restricts the surrogate to interpolation within its trained envelope and assigns all long-horizon forecasting to the physics-based Bayesian filter, which extrapolates through the governing equations rather than through a learned mapping. Section 6.5 shows that the filter performs this role reliably, retaining an 81.5% dispersion reduction over a ten-year unobserved forecast horizon. The answer to the question of what prevents data-driven extrapolation instability in this framework is therefore not that the network is constrained, but that the network is not asked to extrapolate.

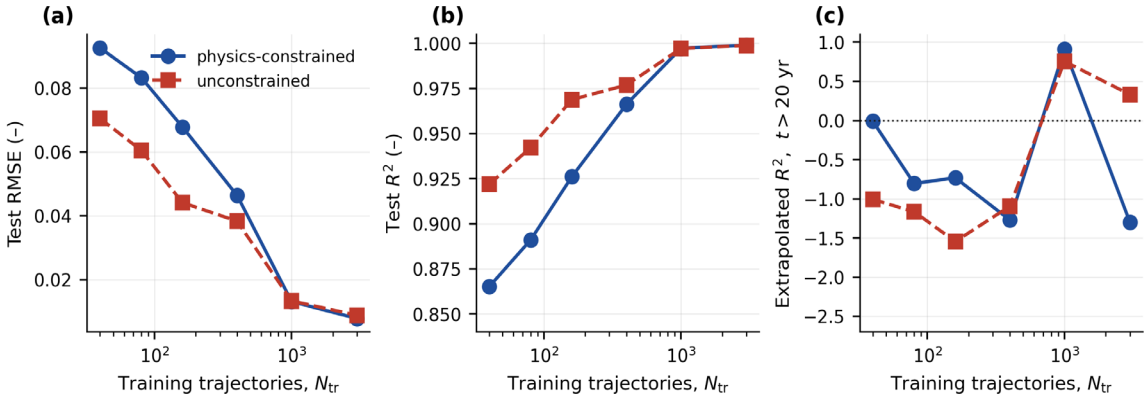


Figure 14. Surrogate performance against training set size. (**a**) Test RMSE. (**b**) Test R^2. (**c**) Coefficient of determination when the model is trained on $t \leq 12$ years and evaluated on $t > 20$ years; values below zero indicate performance worse than the test-set mean.

6.9. Comparison with Deterministic Assessment

Table 14 compares the framework against a deterministic assessment of the same component, in which all inputs are fixed at their mean values.

Table 14. Deterministic assessment against the proposed framework.

Aspect	Deterministic Assessment	Proposed Framework
End-of-life degradation state	0.784 (single value)	0.636 ± 0.300, 90% interval 0.126–0.990
Reliability information	Utilisation ratio only	$\beta(t)$, $P_{f,\beta(t)}$, $P_{f,MC(t)}$, $R(t)$
Year target reliability breached	Not identifiable	17.8 years, shifting to 10.8 years at $\sigma_E = 2.0$
Corrosion–fatigue coupling	Sequential (thickness then fatigue)	Bidirectional, Equations (2), (5) and (6)
Environmental variability	Characteristic value	Full distribution, Equation (18)
Use of inspection data	Requalification at fixed intervals	Continuous assimilation, 81.5% dispersion reduction
Instability detection	Not available	IDII with calibrated thresholds, Table 11
Sensitivity attribution	Not available	Sobol' indices, Table 9

Evaluating the model at mean inputs gives an end-of-life degradation state of 0.784, against a probabilistic mean of 0.636. The deterministic figure is the more severe of the two, because the bounded damage components of Equation (1) respond nonlinearly to their inputs and the expectation of the composite is not the composite of the expectations. A reader might therefore conclude that deterministic assessment is safely conservative.

That conclusion does not hold. The deterministic calculation returns one number and no distribution. It cannot state that the 95th percentile outcome is 0.990, that the target reliability index is breached at 17.8 years, that a fourfold increase in environmental variability would advance that date to 10.8 years, or that assimilating routine inspection data would tighten the forecast by four fifths. Conservatism in the central estimate is not a substitute for knowing the spread, and it is the spread that governs when to inspect and whether to extend life.

7. Implications, Scalability and Limitations

7.1. From Condition Monitoring to Probabilistic Decision Support

Most offshore monitoring systems in service report condition. The framework developed here reports condition, the confidence in that condition, the rate at which confidence is being lost, and the date on which a stated reliability target will be breached. That last output is what converts monitoring into decision support, because it is expressed in the units in which integrity decisions are actually made. The outputs are decision *support* rather than decision *rules*: for the demonstration configuration studied here they establish that the framework produces internally consistent, auditable and reproducible indicators, and Section 4.1 sets out the asset-specific calibration and independent field validation that would be required before those indicators could be relied upon for an operational integrity decision on a specific structure.

7.2. Translating Probabilistic Output into Maintenance Decisions

The link between the probabilistic output and an inspection decision is made explicit through the IDII regime bands of Table 11, which define a monitoring policy: while *IDII* remains below 0.35 the baseline inspection interval is retained; on entering the transitional band the interval is shortened and corrosion protection is reviewed; on entering the unstable band a detailed inspection is triggered and intervention is planned. Because Table 11 also carries the reliability index corresponding to each band, the policy is auditable against the reliability target that governs the asset.

A quantitative example makes the value concrete. For the reference component under baseline conditions, the target reliability is breached at 17.8 years. Under a fixed five-yearly inspection regime the breach would be detected at the year-20 inspection, 2.2 years after it occurred. Monitoring IDII against the transitional threshold of 0.35, which is reached at 10.7 years, gives seven years of warning. The value of the framework in operational terms is that lead time, and Section 6.7 shows that it is robust: the lead time lies between 6.6 and 8.1 years under every one of the 21 composite-index weightings examined. The absolute years quoted belong to the deliberately aggressive demonstration case of Section 4.2; for a code-compliant asset in normal service the same policy would issue its triggers later, and the lead time would have to be re-established from that asset's own calibrated model.

7.3. Climate Resilience

Climate adaptation has become a first-order concern for infrastructure operators, and digital twins have been proposed as the mechanism through which climate information enters operational decision-making [2,51,52], including exposure-state formulations that carry climate stress explicitly into the twin's state vector [53]. Section 6.3 established that environmental

variability degrades reliability even when mean severity is unchanged, and that mean-based assessment can invert the apparent ranking of scenarios. For climate adaptation this has a specific implication: downscaled climate projections used for offshore integrity assessment must preserve the variance and the tail structure of the metocean distribution, not merely the shift in central tendency. A projection that reproduces mean wave height correctly but understates the change in its dispersion will understate the loss of service life.

7.4. Scalability to Full Structural Systems

The reference component is a single member. Real assets are systems of hundreds of members with multiple hot spots, and the computational argument must be made explicitly rather than assumed.

The cost of the framework separates into three parts. The Bayesian assimilation, at 0.62 s per asset for 4000 particles over 20 observations, scales linearly in the number of monitored hot spots and is embarrassingly parallel across them; a 200-hot-spot jacket would require approximately two minutes of single-core time per update cycle, against an annual update interval. This part does not constrain deployment.

The forward model is the binding constraint. In this study the degradation model is a reduced formulation evaluated in closed form, which is why full Monte Carlo is affordable. When the forward model is a nonlinear finite-element system, direct Monte Carlo with 30,000 realisations becomes impractical, and three established strategies apply. Reduced-order models built by proper orthogonal decomposition or reduced-basis methods retain the parametric dependence of the finite-element solution at a fraction of the cost. Surrogate models trained on a designed set of finite-element evaluations—the approach Zhang et al. (2024) applied to offshore wind farm layout optimization [50], where a load surrogate replaced direct simulation inside a genetic algorithm—reduce the per-evaluation cost to microseconds, as Table 5 shows for the surrogate used here. And multi-fidelity or multilevel Monte Carlo schemes allocate most samples to a coarse model and correct with a small number of fine evaluations.

The framework accommodates all three, because the probabilistic and assimilation layers depend on the forward model only through its input–output interface. The important qualification, established in Section 6.8, is that a surrogate substituted in this way inherits a validity envelope, and the sampling design must cover the parameter and time ranges over which it will be queried. System-level effects that a single-member model cannot represent—load redistribution after member failure, correlation of degradation between members exposed to the same environment, and system-versus-component reliability—are outside the present scope and are identified in Section 7.6.

7.5. The Role of Physics-Informed Learning

The evidence in Section 6.8 supports a narrower claim about machine learning in structural digital twins than is usually made. Within the trained envelope the surrogate is excellent and cheap, and physics constraints imposed architecturally guarantee non-negativity and a zero initial state at no accuracy cost in the data-rich regime. Outside that envelope neither the constrained nor the unconstrained network is trustworthy, and the physics constraints tested here did not repair that.

The productive division of labour is therefore to let learned components do what they do well—fast interpolation, pattern recognition, anomaly flagging—and to leave forward projection to the physical model conditioned on data. Bayesian assimilation extrapolates through the governing equations, so its forecasts inherit whatever physical validity those equations possess; a neural network extrapolates through a learned function with no such guarantee. The architecture in Section 3.1 reflects this division deliberately.

7.6. Scope, Assumptions and Limitations

The framework rests on assumptions that bound its applicability, and these are stated plainly.

Validation status. The study is demonstration-oriented. Parameters are drawn from published sources and calibrated against reliability anchors, but the framework has not been validated against a documented in-service asset with traceable inspection records. The results establish internal consistency and correct qualitative behaviour; they do not establish predictive accuracy for a specific structure, and the reference component was configured deliberately so that the reliability target is reached at 17.8 years of a 30-year design life; that crossing year, and the threshold years derived from it, characterise the demonstration case and not typical code-compliant assets. Field validation is the single most important item of further work.

Component scope. A single member with a single hot spot is modelled. System reliability, load redistribution and inter-member degradation correlation are not represented. For redundant structures the component reliability reported here is conservative; for non-redundant (series) configurations the system failure probability is at least that of the governing component, so the component values are unconservative.

Degradation physics. Corrosion is represented by a power law with a lumped severity index, not by a mechanistic electrochemical model, and localised pitting is not distinguished from general wastage. Fatigue uses a single-slope S–N formulation with Miner accumulation, which does not capture load-sequence effects, crack closure or the threshold behaviour that fracture-mechanics models represent. The stress amplification law of Equation (4) is a bounded approximation anchored to an adopted concentration factor at maximum admissible section loss, not a full local stress solution.

Composite index weighting. The weights in Equation (1) are specified rather than derived. Section 6.7 shows that within ±30% perturbations the reliability-target year moves by at most 1.3 years, that the four stress-test weightings confine it to 15.9–19.3 years, that the ordering of the regime transitions is unchanged in every case and that the warning lead time remains within 6.6–8.1 years; the decision output is therefore robust to the assignment, but the numerical thresholds of Table 11 remain a calibration for the baseline weights and would require recalibration if a different weighting were adopted.

Environmental model. The severity index condenses five metocean variables into one scalar, which preserves the marginal variability that Section 6.3 shows to be decisive but discards directionality, seasonality of the joint structure, and extreme-event clustering. Climate non-stationarity is represented through the variability multiplier rather than through downscaled projections, and because σ_E scales the environment-driven scatter of α_c, $\Delta\sigma_0$ and S together with that of Ψ (Section 4.3), the service-life shifts of Section 6.3 are attributable to the environment-driven variability of the problem as a whole rather than to metocean variability in isolation. Ψ is constant within a realisation, so within-life temporal variability of the climate, including storm sequences, is not resolved.

Assimilation assumptions. Measurement errors are assumed independent and Gaussian with known variance. Systematic inspection bias, missed detection and operator-dependent sizing error are not modelled, and all would degrade the reported uncertainty reduction. The 81.5% figure should be read as an upper bound achievable under well-characterised measurement conditions.

Surrogate validity. As established in Section 6.8, the surrogate is valid only within its training envelope. Any

deployment must either retrain as the envelope shifts or fall back to the physical model.

8. Conclusions

This paper has developed and reported in full a physics-informed probabilistic digital twin for offshore infrastructure subject to coupled corrosion–fatigue degradation under environmental uncertainty. The following conclusions are supported by the results.

A bidirectionally coupled degradation model, in which corrosion amplifies the fatigue-driving stress and fatigue damage accelerates corrosion, produces materially different lifecycle behaviour from sequential treatment. Suppressing the fatigue-to-corrosion feedback delayed the breach of the target reliability index by 2.7 years (from 17.8 to 20.5 years) for the reference component, even though the interaction component that records the coupling carries the smallest weight in the composite index.

Regularising the stress-amplification law removes the singularity present in conventional thin-wall formulations while preserving $K_\sigma(0) = 1$ and the adopted concentration factor at maximum section loss, so reliability estimates remain finite and stable as the structure approaches its limit state.

The two-moment failure probability was verified against the empirical failure probability counted from the 30,000 Monte Carlo limit-state samples: the absolute difference is at most 2.2×10^{-4}, the relative difference is below 1% from year 10 onward, and the two-moment value lies within the Wilson 95% interval at every evaluation year, so the Cornell approximation is adequate for this limit state despite the skewness of the degradation state.

The Infrastructure Degradation Instability Index is bounded, dimensionless and free of singularities, and admits a single consistent definition across text, tables and figures. Its regime thresholds of 0.15, 0.35, 0.50 and 0.65 correspond to reliability indices of 3.06, 2.26, 1.71 and 1.05 for the reference component, giving interpretable inspection triggers with a calibrated relationship to the reliability target. A 21-case sensitivity study shows that the composite-index weights move the reliability-target year by at most 1.3 years under ±30% perturbations and confine it to 15.9–19.3 years under four stress-test weightings, leave the ordering of the regime transitions unchanged and keep the warning lead time within 6.6–8.1 years, so the decision logic does not depend materially on the weight assignment.

Environmental variability governs degradation and reliability through dispersion rather than through the mean. Quadrupling the environment-driven variability multiplier at constant mean severity reduced the mean end-of-life degradation state by 34% while increasing its dispersion by 73% and advancing the breach of the target reliability index from 21.5 to 10.8 years. Assessments based on mean environmental conditions can therefore rank scenarios incorrectly, and climate projections used for offshore integrity must preserve variance and tail structure.

Global variance decomposition attributes 86.8% of total-order variance in end-of-life degradation to environmental severity, but only 45.4% of the variance in the limit-state margin, with demand and capacity model uncertainty contributing a further 47.4%. Characterisation effort should be allocated accordingly, and differently depending on whether damage or reliability is the quantity of interest.

Sequential Bayesian assimilation of annual thickness and stress-range measurements reduced the predictive dispersion of the 30-year degradation state by 81.5% across 150 synthetic assets, with the final ten years forecast without observation. Proper resampling and regularisation were essential: batch importance sampling over the same data collapsed to an effective sample size of 20 and produced untrustworthy posteriors.

A physics-constrained neural surrogate reproduces degradation trajectories with RMSE 0.0078 and $R^2 = 0.999$ within its training envelope and guarantees non-negativity and a zero initial state by construction, but neither it nor an unconstrained equivalent extrapolates reliably beyond the trained time horizon. Long-horizon forecasting is therefore assigned to the physics-based Bayesian filter, and the surrogate is confined to interpolation. This is a design decision rather than a limitation, and it is the mechanism by which the framework avoids data-driven extrapolation instability.

The reference component was deliberately configured as an aggressive demonstration case in which the reliability target is reached at 17.8 years; the framework is accordingly presented as probabilistic decision support whose numerical triggers must be recalibrated and validated per asset before operational use. Future work should prioritise validation against instrumented in-service assets with traceable inspection histories, extension from component to system reliability with load redistribution and inter-member correlation, incorporation of downscaled climate projections that preserve variance structure, and treatment of inspection bias and probability of detection within the assimilation likelihood.

Author Contributions: Conceptualization, S.R.S.; methodology, S.R.S.; software, S.R.S.; validation, S.R.S., M.A.G. and N.B.; formal analysis, S.R.S.; investigation, S.R.S., M.A.G., M.A.A.K. and S.P.; data curation, S.R.S.; writing — original draft preparation, S.R.S.; writing—review and editing, S.R.S., M.A.G., N.B., M.A.A.K. and S.P.; visualization, S.R.S.; supervision, S.R.S.; project administration, S.R.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Ethical Approval: Not applicable. This study was based on numerical modelling, probabilistic simulations, and synthetic data and did not involve human participants or animals. Therefore, ethical review and approval were not required.

Informed Consent Statement: Not applicable. This study did not involve human participants, patients, or identifiable personal data.

Data Availability Statement: All model parameters, prior distributions, calibration constants, filter settings and surrogate hyperparameters required to reproduce the results are given in Tables 1–5 and the weight sets of the sensitivity study in Table 12. The Monte Carlo simulations use 30,000 realisations on a quarterly grid over 30 years; the empirical failure probability is counted from the same realisations at years 0, 5, 10, 15, 20, 25 and 30 with Wilson 95% intervals; the sensitivity analysis uses a Saltelli base sample of 4096 giving 49,152 evaluations; the particle filter uses 4000 particles with systematic resampling below an effective sample size of 2000 and multiplicative jitter of 0.03; the 21 weight sets of Table 12 are evaluated from the component means and covariances of the same baseline realisations; the surrogate is a 6–48–48–1 network with tanh hidden activations and a softplus residual output, trained for 400 epochs with Adam at a learning rate of 4×10^{-3} and batch size 128. Random seeds are fixed throughout.

Declaration of the Use of Generative AI and AI-Assisted Technologies: During the preparation of this manuscript, the authors used generative AI-based tools for language editing, grammar checking, and improving the clarity and readability of selected sections of the manuscript. The use of such tools was limited to editorial assistance and did not involve the generation of scientific concepts, research methodology, data, analyses, interpretations, or conclusions. All scientific content, results, and conclusions were independently developed, verified, and critically reviewed by the authors. The authors take full responsibility for the content of the manuscript.

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

Nomenclature

Symbol	Definition
$D(t)$	Composite degradation state (integrated degradation index), $0 \le D \le 1$
$D_c^*,\ D_f^*,\ D_e^*,\ D_i^*$	Normalised (capped) corrosion, fatigue, environmental and interaction damage components; D_f denotes the uncapped Miner sum of Equation (6)
$d_c(t)$	Corrosion penetration depth (mm)
$t_0,\ t_{ca}$	Initial wall thickness and corrosion allowance (mm)
$\alpha_c,\ n_c$	Corrosion rate coefficient and time exponent
t_i	Corrosion initiation time (coating and cathodic-protection life)
$K_\sigma(t)$	Stress amplification factor due to section loss
$\theta(t)$	Section-loss fraction, d_c/t_0
$\Delta\sigma_0,\ \Delta\sigma_{eff}$	Nominal and effective hot-spot stress range (MPa)
$\bar{A}, \log_{10}\bar{A}, m$	S–N curve intercept, base-10 logarithm of the intercept, and S–N curve slope, respectively.
ν	Stress cycle rate (cycles per year)
Ψ	Normalised environmental severity index
$g(\mathbf{X}, t)$	Limit-state function
$\beta(t)$	Second-moment reliability index
$P_f(t),\ R(t)$	Probability of failure and reliability
$P_{f,\beta}(t),\ P_{f,MC}(t)$	Failure probability from the two-moment index, $\Phi(-\beta)$, and the empirical Monte Carlo estimate N_f/N_{MC}
$N_f,\ N_{MC}$	Number of failed realisations and total number of Monte Carlo realisations
$w_c,\ w_f,\ w_e,\ w_i$	Composite-index weights of the corrosion, fatigue, environmental-exposure and interaction components
$\Lambda(t)$	Instability driver
$\hat{\beta}(t)$	Reliability margin normalised by the target index
$IDII(t)$	Infrastructure Degradation Instability Index
σ_E	Dimensionless environmental-variability multiplier that scales the baseline coefficients of variation of Ψ, α_c, $\Delta\sigma_0$ and S; σ_E = 1 is the baseline case; σ_E is not a standard deviation
$S_i,\ S_{Ti}$	First-order and total-order Sobol' indices
$R_0,\ S,\ s_u$	Normalised initial capacity, normalised demand and utilisation ratio, Equation (9)
κ	Capacity transfer coefficient, Equation (9)
$c_\alpha,\ c_\sigma$	Calibrated multipliers on corrosion penetration, Equation (2), and nominal stress range, Equation (5)
$a_{fc},\ K_{\sigma,max},\ c_K$	Fatigue-to-corrosion coupling coefficient and stress-amplification limit and shape constant, Equations (2) and (4)
$T_d,\ D_{crit},\ \sigma_{D,ref}$	Design life, critical degradation level and reference dispersion, Equation (15)
β_{target}	Target reliability index, Equation (16)
$\varsigma_{UT},\ \varsigma_{SG}$	Measurement standard deviations of the ultrasonic thickness and strain-gauge stress-range observations, Equation (20)
$N_p,\ N_{eff}$	Particle population and effective sample size of the sequential filter, Equations (21) and (22)
$\dot{D}(t),\ \dot{D}_{ref}$	Rate of the mean degradation state and its reference value, Equation (15)
$\mu_g,\ \sigma_g$	Mean and standard deviation of the limit-state function, Equation (10)
$D_{phys},\ \mathcal{N}_\Phi$	Known-physics term and neural residual of the surrogate, Equation (23)
z	Standard-normal quantile of the Wilson interval (1.96), Equation (14)

References

[1] Yang, Y.; Chen, C.; Zhuang, Y.; Suo, Z. Reviewing the progress of corrosion fatigue research on marine structures. *Front. Mater.* **2024**, *11*, 1399292. doi: 10.3389/fmats.2024.1399292.

[2] Dale, K.I.; Pope, E.C.D.; Hopkinson, A.R.; McCaie, T.; Lowe, J.A. Environment-aware digital twins: Incorporating weather and climate information to support risk-based decision-making. *Artif. Intell. Earth Syst.* **2023**, *2*(4), e230023. doi: 10.1175/AIES-D-23-0023.1.

[3] Riffat, J.; Samaei, S.R.; Pastakkaya, B. Advances, applications, and future directions of structural health monitoring in civil infrastructure: A comprehensive review. *Energy Catal.* **2025**, *1*, 127–140. doi: 10.61552/EC.2025.010.

[4] Shittu, A.A.; Mehmanparast, A.; Shafiee, M.; Kolios, A.; Hart, P.; Pilario, K. Structural reliability assessment of offshore wind turbine support structures subjected to pitting corrosion-fatigue: A damage tolerance modelling approach. *Wind Energy* **2020**, *23*(11), 2004–2026. doi: 10.1002/we.2542.

[5] Wang, L.; Kolios, A.; Liu, X.; Venetsanos, D.; Cai, R. Reliability of offshore wind turbine support structures: A state-of-the-art review. *Renew. Sustain. Energy Rev.* **2022**, *161*, 112250. doi: 10.1016/j.rser.2022.112250.

[6] Adedipe, O.; Brennan, F.; Kolios, A. Review of corrosion fatigue in offshore structures: Present status and challenges in the offshore wind sector. *Renew. Sustain. Energy Rev.* **2016**, *61*, 141–154. doi: 10.1016/j.rser.2016.02.017.

[7] Shamir, M.; Braithwaite, J.; Mehmanparast, A. Fatigue life assessment of offshore wind support structures in the presence of corrosion pits. *Mar. Struct.* **2023**, *92*, 103505. doi: 10.1016/j.marstruc.2023.103505.

[8] Larrosa, N.O.; Akid, R.; Ainsworth, R.A. Corrosion-fatigue: A review of damage tolerance models. *Int. Mater. Rev.* **2018**, *63*(5), 283–308. doi: 10.1080/09506608.2017.1375644.

[9] Grieves, M.; Vickers, J. Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In Transdisciplinary Perspectives on Complex Systems, Kahlen, F.-J., Flumerfelt, S., Alves, A., Eds.; Springer: Cham, Switzerland, 2017; pp. 85–113. doi: 10.1007/978-3-319-38756-7_4.

[10] Glaessgen, E.H.; Stargel, D.S. The digital twin paradigm for future NASA and U.S. Air Force vehicles. In Proceedings of the 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference, Honolulu, HI, USA, 2012; AIAA 2012-1818. doi: 10.2514/6.2012-1818.

[11] Tao, F.; Cheng, J.; Qi, Q.; Zhang, M.; Zhang, H.; Sui, F. Digital twin-driven product design, manufacturing and service with big data. *Int. J. Adv. Manuf. Technol.* **2018**, *94*, 3563–3576. doi:

10.1007/s00170-017-0233-1.
[12] Kapteyn, M.G.; Pretorius, J.V.R.; Willcox, K.E. A probabilistic graphical model foundation for enabling predictive digital twins at scale. *Nat. Comput. Sci.* **2021**, *1*(5), 337–347. doi: 10.1038/s43588-021-00069-0.
[13] Mousavi, Z.; Varahram, S.; Ettefagh, M.M.; Sadeghi, M.H.; Feng, W.-Q.; Bayat, M. A digital twin-based framework for damage detection of a floating wind turbine structure under various loading conditions based on deep learning approach. *Ocean Eng.* **2024**, *292*, 116563. doi: 10.1016/j.oceaneng.2023.116563.
[14] Pastor-Sanchez, A.; Garcia-Espinosa, J.; Di Capua, D.; Servan-Camas, B.; Berdugo-Parada, I. Real-time digital twin for structural health monitoring of floating offshore wind turbines. *J. Mar. Sci. Eng.* **2025**, *13*(10), 1953. doi: 10.3390/jmse13101953.
[15] Rølvåg, T.; Stranden, Ø. Digital twin based structural health monitoring of offshore crane. In Proceedings of the ASME 2022 41st International Conference on Ocean, Offshore and Arctic Engineering, Hamburg, Germany, 5–10 June 2022. doi: 10.1115/OMAE2022-85896.
[16] Mauro, F.; Kana, A.A. Digital twin for ship life-cycle: A critical systematic review. *Ocean Eng.* **2023**, *269*, 113479. doi: 10.1016/j.oceaneng.2022.113479.
[17] Jorgensen, J.; Hodkiewicz, M.; Cripps, E.; Hassan, G.M. Requirements for the application of the digital twin paradigm to offshore wind turbine structures for uncertain fatigue analysis. *Comput. Ind.* **2023**, *145*, 103806. doi: 10.1016/j.compind.2022.103806.
[18] Ramezani, M.; Choe, D.E.; Heydarpour, K.; Koo, B. Uncertainty models for the structural design of floating offshore wind turbines: A review. *Renew. Sustain. Energy Rev.* **2023**, *185*, 113610. doi: 10.1016/j.rser.2023.113610.
[19] Karniadakis, G.E.; Kevrekidis, I.G.; Lu, L.; Perdikaris, P.; Wang, S.; Yang, L. Physics-informed machine learning. *Nat. Rev. Phys.* **2021**, *3*(6), 422–440. doi: 10.1038/s42254-021-00314-5.
[20] Tao, F.; Sui, F.; Liu, A.; Qi, Q.; Zhang, M.; Song, B.; et al. Digital twin-driven product design framework. *Int. J. Prod. Res.* **2019**, *57*(12), 3935–3953. doi: 10.1080/00207543.2018.1443229.
[21] Li, S.; Brennan, F. Digital twin enabled structural integrity management: Critical review and framework development. *Proc. Inst. Mech. Eng. Part M J. Eng. Marit. Environ.* **2024**, *238*(4), 707–727. doi: 10.1177/14750902241227254.
[22] Li, S.; Brennan, F. Implementation of digital twin-enabled virtually monitored data in inspection planning. *Appl. Ocean Res.* **2024**, *144*, 103903. doi: 10.1016/j.apor.2024.103903.
[23] Sakr, M.; Sadhu, A. Recent progress and future outlook of digital twins in structural health monitoring of civil infrastructure. Smart Mater. Struct. 2024, 33(3), 033001. doi: 10.1088/1361-665X/ad2bd7.
[24] Lai, X.; Yang, L.; He, X.; Pang, Y.; Song, X.; Sun, W. Digital twin-based structural health monitoring by combining measurement and computational data: An aircraft wing example. *J. Manuf. Syst.* **2023**, *69*, 76–90. doi: 10.1016/j.jmsy.2023.06.006.
[25] Chen, C.; Fu, H.; Zheng, Y.; Tao, F.; Liu, Y. The advance of digital twin for predictive maintenance: The role and function of machine learning. *J. Manuf. Syst.* **2023**, *71*, 581–594. doi: 10.1016/j.jmsy.2023.10.010.
[26] Samaei, S.R.; Riffat, J. Intelligent structural health monitoring of jack-up platform legs using high-density sensor networks, real-time digital twins, and machine learning-based damage detection. *Future Cities Environ.* **2025**, *11*, 31. doi: 10.70917/fce-2025-031.
[27] Samaei, S.R.; Riffat, J. Field-validated digital twin–enabled structural health monitoring for offshore high-pile wharves. *Future Cities Environ.* **2026**, *12*, 22. doi: 10.70917/fce-2026-010.
[28] Riffat, J.; Ahadpour Doudran, H.; Samaei, S.R. AI-driven digital twin for uncertainty-aware structural health monitoring of offshore wind turbines considering biofouling effects and reliability prediction. *Green Technol. Innov.* **2025**, *1*, 025330013. doi: 10.36922/GTI025330013.
[29] Riffat, J.; Ahadpour Doudran, H.; Samaei, S.R. AI-enhanced damage detection in jack-up rig legs using an improved modal strain energy index: A numerical, experimental, and digital twin-based approach. *Green Technol. Innov.* **2025**, *1*, 025340014. doi: 10.36922/GTI025340014.
[30] Riffat, J.; Nazari, K.; Samaei, S.R. AI-driven digital twin framework for adaptive and real-time structural health monitoring of offshore marine structures. *Artif. Intell. Sustain. Cities* **2026**, *1*(1), 105–122. doi: 10.65582/aifsc.2026.006.
[31] Melchers, R.E. Modeling of marine immersion corrosion for mild and low-alloy steels—Part 1: Phenomenological model. *Corrosion* **2003**, *59*(4), 319–334. doi: 10.5006/1.3277564.
[32] Melchers, R.E. Modeling of marine immersion corrosion for mild and low-alloy steels—Part 2: Uncertainty estimation. *Corrosion* **2003**, *59*(4), 335–344. doi: 10.5006/1.3277565.
[33] Miner, M.A. Cumulative damage in fatigue. *J. Appl. Mech.* **1945**, *12*(3), A159–A164. doi: 10.1115/1.4009458.
[34] Paris, P.; Erdogan, F. A critical analysis of crack propagation laws. *J. Basic Eng.* **1963**, *85*(4), 528–533. doi: 10.1115/1.3656900.
[35] Shittu, A.A.; Mehmanparast, A.; Hart, P.; Kolios, A. Comparative study between S-N and fracture mechanics approach on reliability assessment of offshore wind turbine jacket foundations. *Reliab. Eng. Syst. Saf.* **2021**, *215*, 107838. doi: 10.1016/j.ress.2021.107838.
[36] Tafazzoli Moghaddam, B.; Hamedany, A.M.; Taylor, J.; Mehmanparast, A.; Brennan, F.; Davies, C.M.; et al. Structural integrity assessment of floating offshore wind turbine support structures. *Ocean Eng.* **2020**, *208*, 107487. doi: 10.1016/j.oceaneng.2020.107487.
[37] Zhu, F.; Yeter, B.; Brennan, F.; Collu, M. Time-domain fatigue reliability analysis for floating offshore wind turbine substructures using coupled nonlinear aero-hydro-servo-elastic simulations. *Eng. Struct.* **2024**, *318*, 118759. doi: 10.1016/j.engstruct.2024.118759.
[38] Yeter, B.; Brennan, F. Probabilistic structural integrity assessment of a floating offshore wind turbine under variable amplitude loading. *Procedia Struct. Integr.* **2024**, *57*, 133–143. doi: 10.1016/j.prostr.2024.03.016.
[39] Partovi-Mehr, N.; DeFrancisci, J.; Minaeijavid, M.; Moaveni, B.; Kuchma, D.; Baxter, C.D.P.; et al. Fatigue analysis of a jacket-supported offshore wind turbine at Block Island Wind Farm. *Sensors* **2024**, *24*(10), 3009. doi: 10.3390/s24103009.
[40] Beck, J.L.; Katafygiotis, L.S. Updating models and their uncertainties. I: Bayesian statistical framework. *J. Eng. Mech.* **1998**, *124*(4), 455–461. doi: 10.1061/(ASCE)0733-9399(1998)124:4(455).
[41] Katafygiotis, L.S.; Beck, J.L. Updating models and their uncertainties. II: model identifiability. *J. Eng. Mech.* **1998**, *124*(4), 463–467. doi: 10.1061/(ASCE)0733-9399(1998)124:4(463).
[42] Vanik, M.W.; Beck, J.L.; Au, S.-K. Bayesian probabilistic approach to structural health monitoring. *J. Eng. Mech.* **2000**, *126*(7), 738–745. doi: 10.1061/(ASCE)0733-9399(2000)126:7(738).
[43] Arulampalam, M.S.; Maskell, S.; Gordon, N.; Clapp, T. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Trans. Signal Process.* **2002**, *50*(2), 174–188. doi: 10.1109/78.978374.
[44] Rocchetta, R.; Broggi, M.; Huchet, Q.; Patelli, E. On-line Bayesian model updating for structural health monitoring. Mech. Syst. *Signal Process.* **2018**, *103*, 174–195. doi: 10.1016/j.ymssp.2017.10.015.
[45] Torzoni, M.; Manzoni, A.; Mariani, S. Enhancing Bayesian

model updating in structural health monitoring via learnable mappings. *Eng. Comput.* **2025**, *41*, 4953–4975. doi: 10.1007/s00366-025-02224-x.

[46] Wang, M.; Wang, C.; Hnydiuk-Stefan, A.; Feng, S.; Atilla, I.; Li, Z. Recent progress on reliability analysis of offshore wind turbine support structures considering digital twin solutions. *Ocean Eng.* **2021**, *232*, 109168. doi: 10.1016/j.oceaneng.2021.109168.

[47] 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. doi: 10.1016/j.jcp.2018.10.045.

[48] Viana, F.A.C.; Subramaniyan, A.K. A survey of Bayesian calibration and physics-informed neural networks in scientific modeling. *Arch. Comput. Methods Eng.* **2021**, *28*, 3801–3830. doi: 10.1007/s11831-021-09539-0.

[49] Dourado, A.; Viana, F.A.C. Physics-informed neural networks for missing physics estimation in cumulative damage models: A case study in corrosion fatigue. *J. Comput. Inf. Sci. Eng.* **2020**, *20*(6), 061007. doi: 10.1115/1.4047173.

[50] Zhang, X.; Wang, Q.; Ye, S.; Luo, K.; Fan, J. Efficient layout optimization of offshore wind farm based on load surrogate model and genetic algorithm. *Energy* **2024**, *309*, 133106. doi: 10.1016/j.energy.2024.133106.

[51] Riaz, K.; McAfee, M.; Gharbia, S.S. Management of climate resilience: Exploring the potential of digital twin technology, 3D city modelling, and early warning systems for climate-resilient cities. *Sensors* **2023**, *23*(5), 2659. doi: 10.3390/s23052659.

[52] Ayadi, R.; Forouheshfar, Y.; Moghadas, O. Enhancing system resilience to climate change through artificial intelligence: A systematic literature review. *Front. Clim.* **2025**, *7*, 1585331. doi: 10.3389/fclim.2025.1585331.

[53] Samaei, S.R.; Riffat, J. Health-aware digital twins for building control under climate extremes: An exposure-state control framework. *Res. Rev. Sustain.* **2026**, *2*(1), 179–203. doi: 10.65582/rrs.2026.011.



Copyright: © 2026 by the author(s). *Net Zero* published by CoSci Press. This is an open access publication under the Creative Commons Attribution 4.0 International License (CC BY 4.0).