

A Sensor-Integrated, Explainable Machine-Learning Digital-Twin Framework for Field-Validated Service-Life Prediction of Cementitious Materials under Coupled Chloride–Sulfate Carbonation Degradation

Devesh Ojha

dojha@lko.amity.edu

Amity University

Pawan Singh

Amity University

Perspective

Keywords: cementitious materials, corrosion, degradation mechanisms, service life prediction, digital twin, explainable machine learning, chloride-sulfate-carbonation coupling, durability, Bayesian updating, physics-informed neural networks

Posted Date: October 12th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11026995/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License. [Read Full License](#)

Additional Declarations: No competing interests reported.

Abstract

Reinforced cementitious infrastructure is subjected to all of these mechanisms – chloride ingress, sulfate attack, carbonation, and freeze-thaw or alkali silica cycling – during its service life, but most of the available service-life models treat these mechanisms separately. There are a number of coupled mechanisms that are not currently covered by the industry standard approaches such as classical deterministic Fick's second law, or semi-probabilistic approaches such as Life-365, DuraCrete and the fib Model Code for Service Life Design, for which it is not yet possible to find broadly accepted models. At the same time, a large number of machine-learning (ML) and deep-learning (DL) studies have developed to be able to predict short-term performance with higher accuracy, but as recently summarized in the RILEM Technical Committee 315-DCS and related reviews, most of the studies are based on laboratory generated data, with only limited field validation, few open datasets and very poor model interpretability. After synthesising available evidence on degradation mechanisms, microstructural characterisation, supplementary cementitious materials, self-healing technologies, classical service-life models, and ML/digital-twin structural health monitoring, the present authors conclude that a specific research gap, identified in the literature, lies in the absence of an explainable, sensor-integrated, digital-twin framework that couples multi-hazard mechanistic degradation models with field-updatable machine learning for cementitious materials. This updated version clearly aligns the proposed contribution with six more recent digital-twin and ML-durability frameworks (Section 3.1); develops the governing mathematical formulations for each layer (Section 5); develops a computational workflow and synchronisation algorithm (Section 5.1); expands the explainable-AI and machine-learning methodology (Sections 5.5–5.6); adds a practical implementation and uncertainty-propagation treatment (Section 5.7 and Section 8). A five-layer conceptual framework is suggested that merges embedded electrochemical and optical sensing with a coupled chloride-sulfate-carbonation reactive-transport core, a physics-informed explainable ML layer, Bayesian updating of the digital twin, and a decision support output that is interpretable. A validation strategy, based on existing accelerated-exposure and field-monitoring datasets, is outlined and the framework is illustrated with literature-derived parameter ranges. The paper is written in a concept and method-oriented manner and not to report new laboratory results.

1. Introduction

The durability of the reinforced steel embedded in the concrete is the basis for the safety, serviceability and life-cycle cost of most of the civil infrastructure worldwide which is constructed from concrete and other cementitious composites. As soon as cracking or excess porosity occurs, aggressive agents like chloride and sulfate ions and dissolved carbon dioxide enter the matrix and depassivate the reinforcement, leading to structural deterioration in many cases, sometimes within a few months of construction. [4] The ability to accurately estimate the remaining service life of such structures is therefore key to the effective inspection, maintenance, and repair planning process [1, 3].

Since then, the field has evolved from basic “rules of thumb” empirical models developed by Tuutti's two-stage initiation-propagation model of reinforcement corrosion [65] to deterministic diffusion-based models developed by Bazant [67], Saetta et al. [68], Tang and Nilsson [69] and Andrade [70]; to fully probabilistic performance-based models like DuraCrete and the fib Model Code for Service Life Design; and, most recently, to data-driven machine-learning models [1, 3, 5, 20]. As this review shows, there is no widely validated model that couples multiple simultaneous degradation processes, that can be field-tested in a real-world setting with sensors, that has an easily understood structure that can be used in engineering decisions, and that can accommodate new monitoring data to improve predictions. This claim is not made on its own but is formally tested against six recent digital-twin and ML-durability frameworks in Section 3.1. Each of these four capacities is explored through a review of the literature individually, it is shown that the intersection of these four is essentially empty, and a conceptual framework, along with a computational workflow, governing mathematical formulation, an illustrative demonstration and a concrete validation strategy, is introduced to start filling that gap.

2. Literature Review

2.1 Chloride-induced corrosion and ingress modelling

Chloride-induced corrosion of embedded steel is the most studied durability mechanism in cementitious materials, which has been the dominant mechanism in marine and deicing-salt environments. The physical basis of most of the models which have followed these treatments is the binding isotherm by Tang and Nilsson [69] and the migration-based determination of the diffusion coefficient by Andrade [70] that were developed in the 1970s, and which are now contained in Neville's and Mehta and Monteiro's standard reference texts on concrete durability [71, 72]. Recently, finite element and probabilistic analyses were used to validate the chloride diffusion coefficient and to show that it is highly dependent on the water-to-binder ratio of the mix with reported values between 4.5 and $8.6 \times 10^{-12} \text{ m}^2/\text{s}$, depending on the mix design and curing age [6]. Models incorporating Faraday's electrolysis principle and Fick's diffusion law have been known to improve the reliability by 15–20% from older models like Life-365 if humidity effects and concrete resistivity are considered [2]. Some pore-specific investigations that have used triple component mixes of slag-fly-ash-inhibitor have measured chloride diffusion coefficients as low as $1.14 \times 10^{-12} \text{ m}^2/\text{s}$, and geometry optimisation to minimise stress concentrations [7]. Previous comparison studies of traditional empirical models with ML techniques have consistently demonstrated higher predictive accuracy for ML models on the test sets but failed to validate them independently in the field [3].

2.2 Sulfate attack and coupled chloride-sulfate mechanisms

The concept of external sulfate attack, which occurs because of the formation of ettringite and gypsum, is now known to be uncommon in practice. Under fully coupled, depth-dependent multi-ion reactive conditions, poroelastic theory has been used to derive models showing that exposure of reinforced concrete piles in the splash zone is most vulnerable due to increased binding affinity of aggressive ions in a dynamic wetting condition [8]. Hydro-chemo-mechanical models that account for moisture transport, calcium leaching and crystalline phase changes also confirm that the kinetics of sulfate attack are significantly different from fully saturated laboratory tests under unsaturated and cyclic wet/dry conditions [9]. Experimental studies on concrete exposed to sea water have demonstrated that the slag dosage prolongs the time of competitive binding between chloride ions and sulfate ions; and that hybrid genetic-algorithm-optimized artificial neural networks (ANNs) outperform simplified multi-factor analytical models in capturing this competition [10]. Two major issues discussed in reviews of drying-wetting sulfate attack modelling are the absence of a single chemo-transport-mechanical coupling and the challenge of predicting the evolution of pore structure in comparison with independent experimental data sets [11, 12, 14].

2.3 Carbonation-induced deterioration

A decrease in pore-solution alkalinity below the passivation threshold ($\text{pH} < 9.8$) can be achieved through carbonation, which has been monitored in the following ways: fibre Bragg grating sensors, embedded optical pH/chloride probes and electrochemical impedance-based models [50, 51, 58]. Most subsequent empirical and data-driven carbonation models have been based on the square-root-of-time carbonation-depth model derived from first principles CO_2 diffusion-reaction model [66] by Papadakis, Vayenas and Fardis. Artificial neural networks are applied to separate the effect of exposure humidity, CO_2 concentration, and cement replacement level on the carbonation depth [53, 54] and multi-gene genetic programming and random forest regression are applied to sets of concrete mixtures of approximately 200 mixes with fly ash. One common shortcoming found in the literature related to carbonation is that most models are developed based on short duration accelerated tests and cannot predict the carbonation behavior over long exposure timescales [55, 57].

2.4 Freeze-thaw cycling, alkali-silica reaction, and other physical mechanisms

Freeze-thaw cycling and alkali-silica reaction (ASR) are relatively under-represented in the literature of durability in term of ML. Progressive loss of mechanical properties as a function of mass-loss and phase changes of engineered cementitious composites and chemically bonded phosphate ceramic cements has been determined through cyclic freeze-thaw testing and analyzed by thermogravimetric and X-ray diffract graphic measurements [4, 29]. The fib Model Code for Service Life Design clearly acknowledges that there is currently no widespread acceptance of predictive models for ASR and sulfate attack and only considers carbonation- and chloride-induced corrosion in the probabilistic scope of the code [62].

2.5 Microstructural characterisation techniques

The toolbox of microstructural characterization techniques that are used to understand degradation mechanisms includes scanning electron microscopy with energy-dispersive spectroscopy (SEM-EDS), X-ray diffraction (XRD), mercury intrusion porosimetry (MIP), thermogravimetric analysis (TGA/DTG), Fourier-transform infrared spectroscopy (FTIR), and X-ray micro-computed tomography (micro-CT) [24, 26, 27, 28, 29, 30, 32, 33]. A framework for quantitative SEM-EDS hypermap analysis and a carbonation front imaging method based on micro-Raman/X-CT have significantly enhanced phase-level quantification of hydration and degradation products, but the intrinsic resolution of EDS and sample preparation artefacts are still recognized limitations [26, 27, 30]. Although these characterisation datasets are crucial, they are not commonly standardised or openly released in a format that is directly usable for building ML features [30].

2.6 Supplementary cementitious materials (SCMs) and mitigation strategies

The use of fly ash, ground granulated blast-furnace slag (GGBS), silica fume, metakaolin and rice husk ash (RHA) has been widely reported to decrease permeability and chloride diffusivity, and to increase corrosion resistance, with fly ash being reported to reduce corrosion rates by approximately 25% and GGBS to reduce chloride diffusion by approximately 40% of the values obtained for plain Portland cement mixes in some studies [2, 34, 35]. The results of the electrochemical half-cell potential tests show that the mixes containing the fly-ash show delayed corrosion initiation, and silica-fume and metakaolin ternary mixes have enhanced compressive strength and sulfate resistance [36, 37, 38]. Broadly speaking, however, reviews indicate that SCM is quite mixed and curing dependent and that high-volume replacement can have a negative impact on the early-age strength development, which can make the predictive modelling of its performance more complex [39, 40, 41].

2.7 Self-healing cementitious materials

Examples of self-healing technologies based on bacteria and microcapsules include microbially induced calcium carbonate precipitation (MICP) using *Bacillus subtilis*, *Sporosarcina pasteurii* and related strains, which have been shown to heal cracks up to about 1 mm wide and reduce permeability by an order of magnitude compared to non-healing controls [42, 43, 44, 45, 46, 47, 48, 49]. These technologies offer the ability to autonomously manage the crack ingress pathways that promote chloride/sulfate/carbonation attack by extending the service life of the concrete, however, the long-term performance of these technologies in the field, when coupled with multi-hazard exposures, is still not adequately quantified [43, 47].

2.8 Classical and probabilistic service-life prediction models

Despite their simplicity, deterministic Fickian models are still popular, albeit single-mechanism models. The conceptual basis for the next generation of diffusion- and corrosion-rate-based service-life models, such as Life-365, continues to be the two-stage model of initiation followed by propagation developed by Tuutti [65] and the electrochemical corrosion-rate formulation for steel in concrete sea structures developed by Bazant [67]. Life-365 extended that by adding diffusion coefficients that vary with the SCM, based on Thomas and Bentz's original computational implementation [73] and has been field-validated for chloride ingress in a marine exposure site in Maine, USA [59, 60]. DuraCrete developed a fully probabilistic methodology for performance-based design which was formalised in the fib Model Code for Service Life Design [59, 61, 62, 64] that includes four levels of sophistication, ranging from full probabilistic analysis to satisfy rules. Although they are sophisticated, they are recognized by their own authors as subject to local application, the use of widely accepted degradation models which are not yet available for ASR and sulfate attack, and the use of parameter sets that are seldom changed after calibration [59, 61, 62].

2.9 Machine learning, deep learning, and explainable AI in durability assessment

The RILEM TC 315-DCS review of machine-learning applications in concrete durability 2013–2024 revealed that the use of data collected directly from the field is still limited and there is too few of new generated, publicly available datasets, most studies have used similar datasets obtained from the literature, limiting the possibility of widespread validation [5]. A parallel review of the application of ML and DL in Sustainable Concrete Construction found prediction accuracy of up to 99% on curated datasets, with data quality issues defined as the main challenges, along with the lack of standardised datasets and the lack of interpretability in models and suggests the use of explainable AI (XAI) frameworks and hybrid physics-data models as a solution [20]. The identified key unresolved challenges found in complementary reviews of AI for crack detection and SHM are the lack of

standardized benchmark datasets, lack of real-time deployment validation, and lack of strong generalisation across domains [22, 23].

2.10 Structural health monitoring, embedded sensors, and digital twins

The idea of digital twins has been introduced in the field of mechanical and aerospace structural health monitoring, and is now proposed to civil infrastructure, where it can be divided into three categories: digital model (manual data exchange), digital shadow (one-way automated data flow), and true digital twin (bidirectional automated data flow) [19]. Aleatoric and epistemic uncertainty has been directly utilized in real-time structural monitoring using a Bayesian approach based on neural networks (NNs) [18]. Both kinds of sensors have been used to monitor, in-situ, the various corrosion-relevant parameters, such as the case of 155 concrete cores that were extracted from the Itaipu Dam and used to validate long term carbonation uptake [50, 51, 54, 56, 58]. The literature for digital-twins and SHM so far, however, is overwhelmingly concerned with structural (crack, fatigue, vibration) damage detection, and only a few studies consider coupled multi-hazard degradation processes as a first-class model component, particularly the chemistry-driven degradation processes described in Sections 2.1–2.4 [16, 17, 18, 19, 20, 21, 22, 23].

2.11 Cross-cutting synthesis, conflicting findings, and research trends

The literature is summarised section by section in the following: 2.1–2.10, which are based on the mechanism-by-mechanism and technology-by-technology approach, this subsection is a more analytical approach, a requirement from a review of Q1-standard.

Conflicting findings. There is no consensus in the mechanistic sulfate-attack literature regarding the prediction of degradation rates in the field: saturated laboratory kinetics predict accelerated ingress compared to fully submerged conditions [8]; hydro-chemo-mechanical models of reaction kinetics including cyclic wet-dry moisture transport predict a different sensitivity, with unsaturated cyclic conditions changing the reaction kinetics in different ways than saturated benchmark tests [9]. These can only be reconciled if the reaction rate depends on moisture state instead of a single calibrated diffusion coefficient, which is used in the mechanistic formulation in Section 5.2. There is a second tension between ML accuracy (which can reach 99% accuracy on curated lab data [20]) and ML reliability (which has been highlighted as the key unresolved issue in the wide literature [5, 23] by independent research into the literature). This mismatch between in-sample accuracy and out-of-sample reliability is the driving force behind the proposed field-updatable, uncertainty-aware architecture presented in Section 4.

Research trends. Three trends appear in the Sections 2.1–2.10. First, the move from single-ion Fickian models to the multi-ion, reactive-transport formulations (sections 2.1–2.2) is almost entirely confined to the arena of mechanistic modelling and hasn't yet been linked to data streams from field sensors. Second, the literature of ML/DL is shifting away from 'black-box' regressors to physics-informed and explainable architectures, but this shift has not yet kept up with the mechanistic-coupling trend, as only a minority of reviews (e.g., RILEM TC 315-DCS, 2025 [5]) report interpretability as a standard practice. Third, chemical-driven durability degradation, involving slower diffusion-controlled processes and different modalities of sensor (electrochemical and optical, as opposed to accelerometric) is not as rapidly maturing as mechanical/structural damage (vibration, fatigue, cracking).

Limitations of the existing literature, taken as a whole. There are four common limitations across the ten sub-domains: (i) Validation to laboratory-exposure data is nearly always done, but rarely to field-exposure data; (ii) Data sets are mostly not open or are not standardised, making independent replication difficult; (iii) Coupling between mechanisms are either mechanistic-only (no ML, Sections 2.1–2.2), or only ML (no mechanistic constraint, Section 2.9); and (iv) Interpretability and uncertainty quantification are considered optional add-ons, not core design requirements. These are the four restrictions, not any one of them missing, that constitute the specific research gap formalised in Section 3.

The direction suggested by the reviewed literature for future work. The sources reviewed suggest the most literature-consistent ways forward: Federated, multi-owner field datasets (responding to (ii)); Hybrid mechanistic–ML architectures with explicit physics-based loss terms (responding to (iii)); and Standardised interpretability reporting (alongside accuracy metrics) (responding to (iv)), each of which is further developed in Sections 5 and 9.

3. Research Gap Synthesis and Novelty Justification

Table 1 summarizes the most common degradation mechanisms and their corresponding characterisation methods, while Table 2 outlines the generations of the service-life prediction model, which now feature physics-informed and graph/attention-based architectures and is compared to four capability dimensions that frequently appear throughout the literature summarized above: multi-hazard coupling, field-data validation, interpretability, and ability to update in real-time. Table 3 also compiles the specific restrictions that have been detailed by four recent and independent review articles, along with the year of publication and an indicative TRL.

3.1 Comparative positioning against existing digital-twin and ML-durability frameworks

This novelty claim is directly tested in Table 3a, by comparing the six references most recently and representative of the examined frameworks from literature in the fields of digital-twin, structural health monitoring, and ML-durability, outlined in Sections 2.9–2.10. The five dimensions used in defining the proposed contribution are used to score each framework.

Table 3

a. Direct comparison of the proposed framework against six recent digital-twin, SHM, and ML-durability studies across the five capability dimensions underpinning the novelty claim.

Framework / study	Domain	Mechanistic coupling	Explainable AI	Bayesian / DT updating	Embedded field sensing	Multi-hazard
Digital twin–driven SHM review [19]	Structural (mechanical) SHM	No – data-driven only	Not addressed	Conceptual DT layers only, no Bayesian scheme	Generic (vibration/strain)	Single (mechanical damage)
GAN-based damage detection, Z24 Bridge benchmark [17]	Structural SHM	No	No	No explicit Bayesian updating	Accelerometer arrays	Single (structural)
Bayesian neural network real-time SHM [18]	Structural SHM	No	Partial (uncertainty decomposition only)	Yes – aleatoric/epistemic separation	Generic structural sensors	Single (structural)
AI–SHM convergence review [22]	Structural / general SHM	No	Partial, narrative only	Not systematically addressed	Various (vision, vibration)	Single to dual
Data-intelligence durability review [23]	Durability / damage diagnosis	Partial – reviewed, not integrated	Not addressed	Not addressed	Various	Reviewed separately, not coupled
RILEM TC 315-DCS ML-durability synthesis [5]	Durability ML, cross-mechanism	No – identifies coupling as a gap	Identifies as a minority practice	Not addressed	Reviewed, not integrated	Reviewed separately, not coupled
Proposed framework (this paper)	Chemistry-driven multi-hazard durability	Yes – coupled Cl ⁻ /SO ₄ ²⁻ /CO ₂ reactive-transport core (Section 5.2)	Yes – SHAP-based, engineer-facing (Section 5.5)	Yes – explicit Kalman/Bayesian formulation (Section 5.3)	Electrochemical + optical, embedded (Layer 1)	Coupled chloride-sulfate-carbonation

A novelty claim will be defended by answering four questions directly derived from Table 3a below.

What is the novel, i.e., the change? All five of the columns are contained in a single architecture, with no reviewed framework; a mechanistic multi-hazard reactive-transport core coupled with field-updatable, explainable ML inside a synchronised digital twin. Bayesian updating is found in [18] for the case of structural SHM, mechanistic coupling is found in [8, 9, 10] for SHM with

chloride-sulfate interaction, and three separate reviews [5, 20, 22] call for explainability, but fail to implement it as a durability DT layer. Not any one of these in isolation, but the integration architecture that allows for (simultaneous) constraints between the mechanistic model and the ML correction term in terms of dataflow and update logic (Section 5.1) is indeed novel.

1. What parts are in place? Each of the five components is discussed on its own terms and is referenced to the literature in which it was developed in the context of embedded electrochemical/optical sensing [50, 51, 56, 58]; Fickian/Nernst-Planck and coupled sulfate reactive-transport modelling [6, 8, 9]; SHAP/attention-based explainability in adjacent engineering ML domains [20, 22, 23]; Bayesian/Kalman-type digital-twin state updating [16, 17, 18]; and structural digital-twin architecture patterns [15, 19]. There are no claims here for any of these.
2. Which of the above integrals is the original? The three integration points are: (a) integrate the mechanistic model governing equations (Section 5.2) as a physics-informed loss term into the ML model (Section 5.4), which is not done in parallel in [8] and [20] respectively; (b) route embedded electrochemical/optical sensor streams into the same Bayesian state-update equation (Section 5.3) as the ML correction, so that the digital twin's posterior reflects both physical and data-driven evidence together, whereas the authors of [5, 20, 22] integrate both in parallel but narratively; and (c) require SHAP-based attribution as a necessary output of each prediction cycle (Section 5.5), which is not done architecturally in [5, 20, 22].
3. How are existing frameworks for digital twin not just extendable? The structural DTs, examined in the reviewed works [15, 16, 17, 18, 19], are based on observations of vibrations, strains, or displacements sampled at sub-second to minute time scales and around damage states which evolve on a month-long timescale. Chemistry-driven durability degradation is governed by diffusion-reaction physics, which operates over years to decades, observed in the form of electrochemical and optical modalities with different noise and drift properties (Section 8), and necessitates a mechanistic core (Section 5.2), which is absent in structural SHM. The more defensible framing is a purpose-built durability DT, which could re-use the generic synchronisation logic of structural DTs, rather than an extension of any of the existing structural DTs, would entail replacing the observation model, the state-transition model and the update timescale.

Table 1
Dominant cementitious degradation mechanisms and their characterisation basis

Mechanism	Description / governing process	Key sources
Chloride-induced corrosion	Cl ⁻ ingress through pores/microcracks depassivates rebar; Fickian/Nernst-Planck transport models widely used	[1, 2, 3, 6, 7, 69, 70]
Sulfate attack	Ettringite/gypsum formation causes expansion, cracking; strongly coupled with moisture and chloride co-ions	[8, 9, 10, 11, 12, 13, 14]
Carbonation	CO ₂ ingress lowers pore-solution pH, destroying the passive film; ANN, electrochemical, and $\sqrt{t}$ models used for depth prediction	[50, 51, 52, 53, 54, 57, 58, 66]
Freeze-thaw / ASR	Cyclic ice-crystallisation pressure and alkali-aggregate expansion cause micro-cracking; least represented in current ML datasets	[4, 29, 61, 62]
Combined/coupled attack	Field structures rarely experience a single mechanism; coupled chloride-sulfate-carbonation models remain mechanistically heavy and rarely field-validated	[8, 9, 10, 14]

Table 2
Comparative summary of service-life prediction model generations (expanded with modern architecture classes)

Model class	Strengths	Limitations	Sources
Fickian diffusion (2nd law)	Deterministic, closed-form, easy to calibrate	Single ion only; ignores coupling and binding non-linearity	[1, 2, 6, 69]
Life-365	Widely adopted in North America; accounts for SCM effect on diffusion	Location-specific coefficients; no explicit sulfate/carbonation coupling	[60, 61, 73]
DuraCrete / fib Model Code	Full probabilistic framework; four design levels (probabilistic to deemed-to-satisfy)	ASR and sulfate attack explicitly excluded due to lack of validated models; static once calibrated	[59, 61, 62, 64]
Standalone ML/DL (RF, ANN, XGBoost)	High curve-fitting accuracy on lab datasets; captures non-linear interactions	Poor field generalisation; limited interpretability; scarce open/standardised datasets	[3, 5, 20, 53]
Physics-informed / hybrid ML (incl. PINNs)	Embeds governing equations as soft constraints; improves extrapolation and reduces data requirement	Still largely untested against long-term field monitoring of coupled degradation; PINN loss-weighting is non-trivial to tune (Section 5.4)	[5, 20, 22, 23]
Graph Neural Networks (GNNs)	Naturally represent spatially distributed sensor networks and pore-network/microstructural topology as graphs	Application to durability degradation is nascent; requires graph-structured field data rarely available today	[5, 20, 22]
Transformers / attention-based models	Capture long-range temporal dependence in multi-year monitoring records; attention weights offer a partial interpretability signal	Data-hungry; attention weights are a weaker interpretability guarantee than Shapley-value attribution (Section 5.5)	[20, 22, 23]
Bayesian Networks	Explicit, human-readable probabilistic dependency structure; naturally supports belief updating as new evidence arrives	Structure learning from durability data is under-explored; scales poorly with the number of coupled variables	[5, 18, 23]

Table 3

Convergent gaps reported by independent recent reviews, with publication year and indicative technology-readiness level (TRL)

Review	Year	Scope	Reported gap	TRL (this domain)	Ref.
RILEM TC 315-DCS	2025	ML models for durability across mechanisms	Heavy reliance on lab data; scarce field data; few open datasets	TRL 2–3 (methodology synthesis)	[5]
ML/DL in sustainable concrete review	2025	Mix design, mechanical & durability prediction	Data-quality gaps, no standard datasets, limited interpretability	TRL 3 (lab-validated components)	[20]
AI-driven crack/SHM review	2025	Vision- and sensor-based SHM	No standardised benchmarks; weak cross-domain generalisation	TRL 4 (component validation in relevant environment)	[22]
Data-intelligence durability review	2025	Durability, damage diagnosis, performance prediction	Gap between experimental progress and field deployment	TRL 3–4	[23]

Combined, these four strands of evidence (plus the direct comparison in Table 3a) point in the same direction: (i) mechanistic models that allow coupling of chloride, sulfate, and carbonation attack exist but are computationally heavy, are often single-purpose and are not widely used to validate with long-term field monitoring [8, 9, 10]; (ii) ML/DL durability models achieve high accuracy on laboratory data, but are trained on non-standardised, largely reused datasets with weak field generalisation and limited interpretability [5, 20, 24]; (iii) digital-twin and SHM research has matured for structural (mechanical) damage detection, but has not been systematically extended into chemistry-driven, multi-hazard durability degradation [19, 22, 23]; and (iv) classical probabilistic frameworks explicitly exclude several coupled mechanisms because the corresponding validated models do not exist [62]. The specific actionable research gap that this paper addresses is the lack of an explainable, sensor-integrated digital-twin framework that couples multi-hazard mechanistic degradation models with field-updatable machine-learning models for

cementitious materials, which has been validated in collaboration with long-term field data, and is formally distinguished in Table 3a from nearby existing frameworks, rather than simply asserted as such.

4. Proposed Framework

Based on the gap identified in Section 3, a 5-layer conceptual framework is suggested (Fig. 1). The framework is designed in a modular fashion, where each layer can be developed, calibrated, and tested with data types that are already published in the literature and don't necessarily need a new experimental campaign to test the framework. The five layers described qualitatively below are then operationalised below in Section 5, where the computational workflow, governing equations and machine-learning methodology are developed.

4.1 Layer 1 – Field sensing

The relevant parameters that determine the initiation of corrosion are measured in-situ by embedded electrochemical sensors (half-cell potential, linear polarisation resistance, and resistivity probes), and in-situ by optical sensors (fibre Bragg grating strain/temperature, pH- and chloride-sensitive optodes), with periodic ex-situ sampling of extracted cores for microstructural ground-truthing via SEM/XRD/MIP [26, 50, 51, 56, 58].

4.2 Layer 2 – Coupled mechanistic core

The physically constrained backbone of the model is defined by a multi-ion reactive-transport core, which is based on coupled Fickian/Nernst-Planck diffusion and sulfate-driven ettringite formation kinetics, and carbonation front propagation, a framework already shown for pairwise coupling [8, 9, 10]. The governing equations of this core are derived in full in Section 5.2.

4.3 Layer 3 – Explainable machine-learning layer

The output of Layer 2 is then combined with the sensor stream from Layer 1 and a physics-informed, hybrid ML layer (e.g., gradient-boosted trees or physics-informed neural networks) is trained on the combined output of Layer 2 and the sensor stream from Layer 1, with each prediction decomposed into the contribution of individual exposure and mix-design variables using SHAP-value explainability, thus addressing the interpretability gap discussed in Section 2.9 [20, 22, 23]. The training workflow and explainability formulation are defined in Sections 5.4–5.6, respectively.

4.4 Layer 4 – Digital-twin synchronisation

A Bayesian data assimilation scheme is employed to update the digital twin's internal state by taking account of new sensor readings and explicitly distinguishing between aleatoric (measurement/material variability) and epistemic (model-form) uncertainty, inspired by other approaches used in the wider structural digital-twin literature [16, 17, 18, 19]. To simplify the synchronisation logic in Fig. 1 into a single box, Table 5 also breaks up all of Layer 4 into its eight sub-components, each with a function and representative implementation technology, and Section 5.1 defines the data flow between all the sub-components.

Table 5
Expanded digital-twin architecture: eight constituent sub-components of Layer 4

Sub-component	Function within the digital twin	Representative technology
Physical asset	The instrumented concrete/reinforced-concrete structure being monitored; the source of all ground-truth degradation	In-service bridge deck, marine pile, or de-icing-salt-exposed slab
Data acquisition	Samples embedded sensor channels at a fixed cadence and packages readings with timestamps and sensor IDs	Half-cell potential, LPR, resistivity, FBG, pH/Cl ⁻ optode data loggers
Cloud database	Persists raw and quality-checked time-series data with provenance metadata for later retraining and audit	Time-series database (e.g., structured columnar store) with versioned schema
Edge computing	Performs local pre-processing, outlier screening, and data compression before transmission to reduce bandwidth and latency	On-site microcontroller/edge gateway running lightweight preprocessing
ML engine	Hosts the physics-informed ML layer (Section 5.4); produces point predictions and SHAP attributions on demand	Gradient-boosted trees or PINN, retrained on a scheduled cadence
Synchronisation engine	Executes the Bayesian state-update step (Section 5.3), reconciling the mechanistic-core prior with sensor and ML-layer evidence	Kalman-filter or particle-filter implementation
Feedback loop	Flags divergence between predicted and observed states beyond tolerance, triggering recalibration or maintenance alerts	Threshold-based alerting integrated with the decision layer
Decision layer	Translates the updated posterior remaining-service-life distribution into ranked, interpretable maintenance guidance for engineers	Dashboard reporting RSL credible intervals and top SHAP-ranked drivers (Layer 5)

4.5 Layer 5 – Decision-support output

The framework produces a probabilistic remaining-service-life estimate along with a set of maintainability triggers and an ordered list of the exposure or design variables that are currently most critical to the prediction, which are expected to be easily interpretable by the asset-management engineer and not a data scientist.

5. Computational Workflow, Algorithm, and Mathematical Formulation

In this section, we tackle the methodological question of how these sensor data, and the mechanistic model interact with the ML model and Bayesian updating, respectively, in the original framework presented. The end-to-end computational pipeline and synchronisation algorithm are detailed in Section 5.1, while the mechanical equations, Bayesian updating, and physics-informed ML layer are detailed in Sections 5.2–5.4 respectively. Section 5.5 details the formulation of the explainable AI, while Section 5.6 details the machine-learning workflow. Section 5.7 treats uncertainty propagation explicitly.

5.1 End-to-end computational pipeline and synchronisation algorithm

Figure 1 and Table 5 provide a structural description of the framework's layers; the pipeline below indicates how the data is passed between layers at each synchronisation cycle k (e.g., daily or weekly as per the sensor sampling rate):

Step 1 – Acquisition: Layer 1 sensors transmit raw electrochemical/optical data readings y_k , along with the timestamp and channel ID, to the edge, where it is checked for dropout and gross outliers.

Step 2 – Mechanistic forward pass: The coupled reactive-transport core (Layer 2, Section 5.2) carries out a forward pass between t_{k-1} and t_k with the current state estimate $\hat{x}_{k-1}$ and a prior process-noise covariance, generating the prior predicted state estimate $\hat{x}_{k-1}$ and a prior predicted process-noise covariance.

In Step 3, Physics-informed ML correction: the ML layer (Layer 3, Section 5.4) receives the mechanistic prior and the raw sensor stream as input and outputs a bias-correction term and its own prediction variance, both of which are learned with the physics-informed loss presented in Section 5.4.

In Step 4 – Bayesian assimilation: the synchronisation engine (Section 5.3), the mechanistic-plus-ML prior is used as the prior for the state vector and, similarly to the Kalman filter, is updated with the observed sensors readings y_k to obtain the posterior estimate $\hat{x}_k$ and the updated covariance P_k .

Step 5 – Remaining-service-life (RSL) projection: the posterior state is propagated forward under the mechanistic core until a critical threshold (e.g., critical chloride content at the rebar, or carbonation depth equal to cover depth) is reached, giving a posterior RSL distribution.

At step 6, the current prediction is passed to Layer 5 for engineer-facing reporting and maintenance-trigger evaluation along with SHAP attributions (Section 5.5) and the RSL credible interval.

Steps 1–6 are formalised in pseudocode in Algorithm 1; to allow for different model architectures, the algorithm is deliberately model-agnostic; this means that any admissible mechanistic core (Section 5.2) and any admissible ML architecture (Section 5.4, such as gradient-boosted trees or a PINN) can be replaced without altering the synchronisation logic.

5.2 Mathematical formulation of the coupled mechanistic core

Below are the three most important degradation mechanisms discussed in Sections 2.1–2.3, which are often referred to as transport-controlled. All three have the same reactive-transport structure and are coupled via common state variables - moisture content, temperature, and in the case of chloride-sulfate interaction, competitive ion binding - as discussed in [8, 9] as poroelastic and hydro-chemo-mechanical couplings.

5.2.1 Chloride transport (Fick's second law with binding)

$$\partial C / \partial t = \partial / \partial x [D_{\text{eff}} (\partial C / \partial x)] - \partial C_b / \partial t \quad (1)$$

These are the free chloride concentration (C), the effective diffusion coefficient (D_{eff} , calibrated according to Section 2.1 [6, 7]), and the bound chloride (C_b) concentration, which is related to C by the binding isotherm of Tang and Nilsson [69]. The surface chloride content, C_s , is the boundary condition and the initial condition is $C(x,0) = C_0$.

5.2.2 Nernst–Planck transport (electrochemical migration term)

$$J = -D (\partial C / \partial x) - (zF D C / R T) (\partial \phi / \partial x) \quad (2)$$

The second term (migration) is still included in the framework core as models based purely on diffusion (Fickian) are known to under-predict the ingress once active corrosion currents are established [2].

5.2.3 Sulfate reaction-diffusion (ettringite-forming kinetics)

$$\partial C_{\text{SO4}} / \partial t = \partial / \partial x [D_s (\partial C_{\text{SO4}} / \partial x)] - k_r(\theta_w) C_{\text{SO4}} C_{\text{CA}} \quad (3)$$

where C_{SO4} is the concentration of sulfate ions in the pore-solution, D_s is the diffusion coefficient of sulfate ions, C_{CA} is the concentration of the reactive calcium-aluminate phases available for ettringite formation, and $k_r(\theta_w)$ is a moisture-state-dependent reaction-rate coefficient that is explicitly introduced here, following the cross-cutting synthesis in Section 2.11, to reconcile the conflicting saturated/unsaturated kinetics reported in [8] and [9]. The volume fraction of ettringite is linearly coupled to expansive strain, as in [8, 9].

5.2.4 Carbonation front propagation

$$x_c(t) = k_c \sqrt{t}, \quad k_c = \sqrt{2 D_{\text{CO2}} C_{\text{CO2}} / a} \quad (4)$$

The carbonation model is based on the following equation, which is derived from first principles CO₂ diffusion-reaction theory [66] formulated by Papadakis et al., which relates the depth of carbonation to ambient CO₂ concentration, effective CO₂ diffusion coefficient and the CO₂ buffering (binding) capacity of the cementitious matrix, which depends on the replacement level of SCM (Section 2.6):

The mechanistic core $M(\cdot)$ mentioned in Algorithm (1), Step (2) is comprised of equations (1)–(4) and is solved numerically (finite-difference or finite-element discretisation in x and t) at every synchronisation cycle to yield the previous state $\hat{x}_k^-$.

5.3 Bayesian digital-twin updating formulation

As new sensor evidence is received, Layer 4 (Section 4.4, Table 5) takes this information into account to update the digital twin's belief about the true internal state of the system. For a given prior probability $p(\theta)$ over state/parameter vector θ (e.g., D_{eff} , C_s , k_c) and the likelihood $p(y|\theta)$ of the sensor readings y , the posterior is

$$p(\theta | y) \propto p(y | \theta) p(\theta) \quad (5)$$

This simplifies the standard Kalman recursion (Algorithm 1:) for the linear-Gaussian approximation implemented in the synchronisation engine.

$$\hat{x}_k = \tilde{x}_k^- + K_k (y_k - H \tilde{x}_k^-), K_k = P_k^- H^T (H P_k^- H^T + R_k)^{-1} \quad (6)$$

H is a map from the internal state to the observable sensor quantities and R_k is the sensor-noise covariance (Section 5.7). This formulation is like that presented in [18] for structural digital twins where the state vector is composed of structural damage indices instead of mechanistic-core parameters as in Section 5.2 of this document.

5.4 Physics-informed machine-learning formulation

The ML correction term f_{ML} in Algorithm 1 (Step 3) is trained to minimize a composite loss, which is a combination of data fidelity loss and physical consistency loss:

$$L = L_{data} + \lambda_{phys} L_{phys}, L_{data} = (1/N) \sum_i (\hat{y}_i - y_i)^2 \quad (7)$$

$$L_{phys} = \|\partial \hat{C} / \partial t - \partial / \partial x [D_{eff} (\partial \hat{C} / \partial x)] + \partial C_b / \partial t\|^2 \quad (8)$$

where $\hat{y}_i$ are ML predictions, y_i are observed sensor/lab values, $\hat{C}$ is the network's implicit concentration field, and λ_{phys} is a weighting hyperparameter (tuned as part of the hyperparameter search in Section 5.6) that trades off fit to sparse field data against consistency with the governing transport physics; this is the standard physics-informed neural network (PINN) construction, adapted from the structural digital-twin literature [18] to the chloride-transport Eq. (1) as the exemplar physical constraint, with directly analogous loss terms for Equations (3)–(4).

The remaining service-life projection in Algorithm 1 (Step 5) is obtained by working out the mechanistic core forward from the posterior state $\hat{x}_k$ until the concentration/depth variable is deemed to hit a criterion:

$$t_{RSL} = t^*_{crit} - t_k, \text{ where } \hat{C}(x_{cover}, t^*_{crit}) = C_{crit} \text{ (or } x_c(t^*_{crit}) = x_{cover}) \quad (9)$$

C_{crit} the critical chloride threshold, x_{cover} the cover depth, t^*_{crit} is the solution of a root finding on the forward-projected state trajectory.

5.5 Explainable AI formulation and methodology

Interpretability was noted as a persistent and unaddressed challenge in three separate reviews [5, 20, 22] in Section 2.9. This subsection includes a thorough explanation of why the technique is chosen over alternatives, how feature importance is calculated, and how engineers should use the results to understand the technique.

5.5.1 Candidate explainability techniques

SHAP (SHapley Additive exPlanations): game-theoretic, additive feature-attribution method that provides local accuracy and consistency properties (selected — see 5.5.2).

LIME (Local Interpretable Model-agnostic Explanations): trains a local surrogate linear model around each prediction, which is computationally cheaper than SHAP but not guaranteed to be consistent with Shapley values and may be unstable with resampling.

Attention mechanisms: show learned attention weights in transformer-based architectures (Table 2); have intuitive meaning for sequence data, but attention scores are not a formally guaranteed measure of feature importance.

Integrated gradients: attribute a prediction to input features by computing gradients over a path of inputs from a baseline input; somewhat less naturally applicable to the gradient-boosted-tree variant of f_{ML} than to differentiable networks (like the PINN in Section 5.4).

5.5.2 Why SHAP is selected

SHAP is used as the main explainability layer for the framework for three reasons that align with the interpretability requirements listed in Section 2.9. The first property is local accuracy—LIME and raw attention weights do not ensure this property, which holds for the Shapley-value formulation. The first is local accuracy, meaning that the attributions sum exactly to the model's output, which is not guaranteed by LIME or raw attention weights. Second, SHAP has an efficient implementation for gradient-boosted-tree (GBT) architectures, which Table 2 shows to be the most validated ML class for this domain, and a sampling-based implementation for the PINN variant, so both admissible f_{ML} architectures in Algorithm 1 can share the same explainability layer without changes. Third, it is additive and this property aligns directly with the output demanded by the engineer, which is ranked-driver output, as specified in Layer 5 (Section 4.5).

5.5.3 Computation of feature importance

$$f(x) = \phi_0 + \sum_i \phi_i, \quad \phi_i = \sum_{S \subseteq F \setminus \{i\}} \frac{|S|! (|F| - |S| - 1)!}{|F|!} [f(S \cup \{i\}) - f(S)] \quad (10)$$

where F is the full feature set and the resulting feature vector (e.g., w/b ratio, exposure duration, cover depth, resistivity, cumulative chloride/sulfate/ CO_2 exposure), ϕ_0 is the baseline output (or expected output) of the model and ϕ_i is the Shapley value, i.e., the average marginal contribution of feature i across all feature subsets S . In production, TreeSHAP computes Eq. (10) in polynomial time, while KernelSHAP approximates it by performing weighted linear regression on a set of sampled coalitions for the PINN variant.

5.5.4 How engineers interpret the output

In Layer 5 (Section 4.5), the values of ϕ_i for each prediction cycle are displayed as a horizontal bar chart, where the exposure or design variables that are currently driving the RSL estimate down and up are indicated, and where the numeric RSL credible interval (from Section 5.7) is also shown. This aims at allowing an asset-management engineer without any data-science knowledge to detect the most important impacts of the prediction and to validate this result with physical intuition, directly implementing the interpretability suggestion in [5, 20, 22].

5.6 Machine-learning workflow

Also mentioned in Section 2.9 is that there were no complete ML workflows that were previously defined. This workflow is followed if f_{ML} is trained before deployment in the field, or during the periodic retraining step in Algorithm 1, and the f_{ML} is either a gradient-boosted-tree or a PINN.

Workflow stage	Specification
Training data	Pooled accelerated-exposure datasets (RCPT, salt-ponding, sulfate/carbonation chamber records) and available field-monitoring records identified in Table 4, combined with mechanistic-core outputs (Section 5.2) as physics-consistent synthetic augmentation
Preprocessing	Missing-value flagging (not imputed at this stage), unit harmonisation across sources, IQR-based outlier screening (values beyond 1.5× IQR flagged for manual review rather than automatically discarded)
Feature engineering	Derived variables: w/b ratio, SCM replacement level, cover depth, cumulative chloride/sulfate/CO ₂ exposure, resistivity, exposure-zone classification (submerged/splash/atmospheric), age at first measurement
Normalisation	Per-feature z-score standardisation (fit on the training split only, applied to validation/test splits to avoid leakage)
Train-test split	Structure-level (not record-level) split to prevent leakage across time points from the same structure: 70% structures for training, 15% for validation, 15% held out for test; for field data, a temporal split (train on earlier years, test on later years) is used in addition, following the recommendation in [5] that field validation should be prospective rather than random-split
Hyperparameter optimisation	Bayesian optimisation (or grid search for small parameter spaces) over tree depth/learning rate/number of estimators (gradient-boosted variant) or network width/depth and λ_{phys} (PINN variant, Eq. 7)
Cross-validation	5-fold cross-validation on the training split for hyperparameter selection; leave-one-structure-out cross-validation reported separately as the primary field-generalisation metric
Evaluation metrics	RMSE and MAE on concentration/depth predictions; R ² for goodness of fit; prediction-interval coverage probability (PICP) and calibration curves for the Bayesian/uncertainty output (Section 5.7), reported alongside SHAP summary plots as a standard deliverable

5.7 Uncertainty propagation

Material uncertainty (mix design and workmanship uncertainty), environmental uncertainty (stochastic exposure conditions uncertainty) and model uncertainty (mechanistic and ML model-form uncertainty). The framework deals with these by following the conventional decomposition of aleatoric/epistemic components that are common to structural digital twins [18]:

$$\text{Var}[y] = E_{\theta}[\text{Var}(y|\theta)] + \text{Var}_{\theta}[E(y|\theta)] \quad (11)$$

The first term (aleatoric) accounts for the irreducible variance of sensors and material; it is estimated from multiple sensor channels and the reported laboratory variance. The second term (epistemic) accounts for the uncertainty of the model's form; it can be estimated by ensembling (multiple instances of gradient-boosted-trees) or Monte Carlo dropout (PINN variant). For propagation of material and environmental uncertainty, Monte Carlo sampling of the ranges of input distributions (e.g. D_{eff} , C_s , reported in the literature in Section 2.1) are used through the mechanistic core (Section 5.2), while model uncertainty is propagated via the Kalman covariance P_k in Eq. (6). In the case of the RSL, the reduction in epistemic uncertainty with subsequent field observations is exactly what the Bayesian update in Eq. (6) achieves, so the RSL's width is expected to decrease as the number of field observations increases, although sensor drift (Section 8) needs to be addressed.

6. Illustrative Demonstration

To demonstrate how the mechanistic core (Layer 2, Section 5.2) and comparison logic of the framework would act, the chloride ingress profiles calculated from the classic Fickian solution (Eq. 1) are shown in Fig. 2 for diffusion coefficient and surface chloride content values selected from the range of reported values from the literature [6, 7]. This is an illustrative calculation on literature reported ranges of parameter values and not a new laboratory measurement and is included to show the mechanistic layer of the model, not to report new experimental data.

To compare qualitatively the five model classes with the four capability dimensions, as shown in Table 2, the scores were calculated on a scale of 0–5, based on the strengths and limitations reported in the referenced reviews, but not in a controlled

head-to-head experiment. It is added for visualisation purposes only and not to statistically demonstrate the position of the proposed framework compared to any existing approaches.

6.1 Scoring methodology for Fig. 3

Since a 0–5 qualitative score has no inherent units, earlier the scores in Fig. 3 had been reported without a methodology, which could be subject to a reasonable challenge for their objectivity. The scoring system is as defined below.

A rubric of five points was developed to score each of the four capability dimensions (multi-hazard coupling, field-data validation, interpretability, real-time updating) of each model class in Table 2, based on information found in the literature, with a score of 0 indicating that the capability is not yet described or excluded from the model class and 5 indicating that the capability is standard, field-validated procedure in the reviewed literature for that model class.

Each capability claim from Sections 2 and 3 was independently scored against the rubric, with an average score computed for claims for the same model class and dimension where multiple sources made claims about each class and dimension (e.g., the score for Field-data validation is an average of the lab-only score reported in [3] and [53]). This is a literature derived, rubric-based scoring which is consistent with a structured expert-judgement (Delphi-style) approach, but not an unweighted subjective assignment, and is an explicit limitation (see Section 9). This is also a synthesis of reported evidence, rather than a controlled, blinded benchmarking exercise.

7. Validation Strategy and Data Sources

The framework is proposed at a conceptual level, so Table 4 shows how each layer might be validated using named data types and, where publicly documented, named datasets and benchmarks reported in the literature that were reviewed in Section 2, before a dedicated new field campaign. The validation approach is staged and aims to enable the framework to be tested with existing, citable and verifiable data sources rather than with speculative input or unverifiable input.

Table 4

Candidate validation data sources for each framework layer, with named datasets, indicative sample size, and validation metrics

Data source type / named dataset	Description	Indicative sample size	Validation use	Validation metric
Accelerated lab exposure archives (RCPT, salt-ponding, wet-dry cyclic sulfate, carbonation chamber datasets, e.g. the ~ 200-mixture fly-ash carbonation dataset in [53])	RCPT, salt-ponding, wet-dry cyclic sulfate, carbonation chamber datasets reported in reviewed studies	~ 150–300 specimens per mechanism, per compiled dataset	Cross-validation of the mechanistic core (Layer 2, Section 5.2)	RMSE / R^2 of predicted vs. measured chloride profile or carbonation depth
Long-term field monitoring records (e.g., Itaipu Dam 155-core carbonation record [54]; Maine marine exposure site underpinning Life-365 field calibration [59, 60])	Marine-exposure structures instrumented with embedded chloride/pH/resistivity sensors, or destructively cored at intervals	~ 150–160 cores / multi-decade single-site record	Calibration and drift-correction of Layer 4 Bayesian updating (Section 5.3)	Prediction-interval coverage probability (PICP); posterior covariance shrinkage over time
FHWA Long-Term Bridge Performance (LTBP) Program field dataset	Multi-structure, multi-decade U.S. bridge condition and exposure monitoring programme, publicly documented by the U.S. Federal Highway Administration	Hundreds of instrumented bridge structures nationally	Structure-level, out-of-sample field-generalisation test (leave-one-structure-out, Section 5.6)	Leave-one-structure-out RMSE/MAE; calibration curve for RSL credible intervals
Existing service-life model outputs (Life-365, DuraCrete, fib MC predictions on benchmark structures)	Life-365, DuraCrete and fib MC predictions on benchmark structures	Model-dependent; benchmark-structure count typically 5–20 per comparative study	Baseline comparison for the proposed framework (Fig. 3, Section 6.1)	Relative RMSE reduction vs. baseline model; qualitative capability score (Table 2)
Microstructural characterisation datasets (SEM-EDS, XRD, MIP, micro-CT)	SEM-EDS, XRD, MIP, micro-CT phase and pore-structure data	Sample-dependent; typically, tens of specimens per study	Feature-engineering inputs and physical-plausibility checks for the ML layer (Section 5.6)	Physical-plausibility screening (e.g., predicted C_{CA} consistent with measured phase fractions)

The principles outlined in the reviewed literature for any future implementation are as follows: (i) accelerated laboratory datasets should include exposure duration, mix design and test standard, as recommended by RILEM TC 315-DCS [5]; (ii) field datasets should include sensor calibration procedures and drift-correction methods, as shown in the embedded-sensor literature [56, 58]; (iii) any ML component should include interpretability outputs, such as SHAP values, alongside accuracy metrics, addressing the interpretability gap identified in Section 2.9 [20, 22].

8. Practical Implementation Considerations

If a conceptual framework is to have any use, it must be deployable, and this section deals with the practical constraints that a field deployment of Layers 1–4 would require to satisfy: sensor cost and maintenance, communication protocols, power, long term drift, cyber security, BIM integration and cloud storage.

Consideration	Discussion
Sensor cost	Embedded electrochemical sensors (half-cell, LPR, resistivity) are low-unit-cost relative to optical fibre Bragg grating (FBG) and chloride/pH optodes, which require interrogator hardware; a full Layer-1 deployment on a single structure is expected to be dominated by installation labour and interrogator equipment rather than per-sensor unit cost.
Maintenance	Embedded sensors cannot generally be repaired once cast into cover concrete, so the framework relies on planned redundancy (multiple co-located sensors per zone) and the Bayesian update (Section 5.3) to tolerate individual channel dropout, with periodic ex-situ coring (Section 4.1) providing an independent maintenance-free ground-truth check.
Communication protocols	Low-power wide-area protocols (e.g., LoRaWAN or NB-IoT) suit sparse, infrequent readings from distributed structures, while wired field-bus protocols (e.g., RS-485/Modbus) suit densely instrumented single structures with local edge computing (Table 5); the choice interacts directly with the power budget below.
Power consumption	Continuous electrochemical monitoring is power-intensive relative to the duty-cycled sampling regime (e.g., readings every few hours rather than continuously) that Layer 4's synchronisation cadence (Section 5.1) actually requires, so duty-cycled sampling combined with energy harvesting (solar, or structural vibration harvesting where available) is expected to be sufficient rather than continuous mains power.
Long-term drift	Electrochemical reference electrodes and optical sensors are both known to drift over multi-year deployments; the framework treats drift as an epistemic-uncertainty contribution (Section 5.7) that the periodic ex-situ coring and redundant-channel cross-checks (Table 5, feedback loop) are designed to detect and correct via recalibration.
Cybersecurity	Because Layer 4's synchronisation engine (Section 5.3) automatically updates a structure's official condition assessment, the data-acquisition and cloud-database components (Table 5) require authenticated, encrypted telemetry and access-controlled write permissions, to prevent tampering with data that could otherwise trigger false maintenance decisions or mask genuine deterioration.
BIM integration	Sensor locations, predicted RSL, and SHAP-ranked drivers (Section 5.5) are intended to be exposed as attributes on the corresponding structural elements in an IFC-based BIM model, so that Layer 5's decision-support output (Section 4.5) sits inside the same asset-management environment engineers already use rather than a standalone dashboard.
Cloud storage	Given multi-decade service lives, the cloud database (Table 5) needs an explicit data-retention and schema-versioning policy so that historical sensor streams remain usable for retraining (Section 5.6) even as sensor hardware, ML architectures, and mechanistic-core implementations are replaced over the structure's lifetime.

9. Discussion

The choices of technical building blocks for the framework proposed here are intentionally conservative, in the sense that each layer is based on a mechanism, sensor type or modelling approach that has been demonstrated individually, in the literature surveyed in Section 2. Importantly, it brings these previously separate threads together explicitly, with a goal of filling the convergent gaps identified, independently, by four recent review articles [5, 20, 22, 23] and only formally distinguished from the closest existing frameworks in Table 3a. This integration is supposed to have three benefits compared to the current integration strategies. First, the coupling of the mechanistic core with field sensing directly solves the RILEM TC 315-DCS problem [5] that durability models for ML were not validated with field data frequently. Second, the explicit explainability layer (Section 5.5) is called for by the repeated calls in at least three independent reviews to interpretable and physically constrained ML in this domain [20, 22, 23]. Third, Bayesian digital-twin updating (Section 5.3) takes the predominantly structural-damage oriented digital-twin literature out of the domain of structural damage and into the realm of chemistry-driven durability [16, 17, 18, 19].

There are also some restrictions on the framework. The illustrative results in Section 6 are based on literature-derived parameter ranges and the qualitative capability scores in Fig. 3 are based on reported strengths and limitations and are presented in the context of an explicit rubric, provided in Section 6.1, but should not be interpreted as performance claims; both should be understood as motivating illustrations rather than valid performance claims. The computational cost of coupled multi-hazard mechanistic models are still fairly high, and integration with real-time model updating of the ML layer leads to open research questions on the computational cost and the potential risk of the ML layer overwriting physically implausible mechanistic constraints (which may be overcome using the physics-informed loss term in Eq. (8) but is not guaranteed). Sulfate attack and

alkali-silica reaction, especially, are not yet well understood with well accepted mechanistic models as the fib Model Code identifies as a precondition to probabilistic design [62] and any attempt to implement Layer 2 for these mechanisms will have to rely on the less well accepted models described in Sections 2.2 and 2.4 [8, 9, 10, 11]. Last, and most importantly, an embedded sensor must be reliably installed, and must possess long service lives – in the order of decades – for the service life of any concrete infrastructure, but not be an additional source of ingress pathways as described in the work on sensors [56, 58] and addressed in Section 8.

9.1 Additional open questions: complexity, scalability, and data scarcity

In addition to the above constrictions, four additional questions arise naturally from the workflow in Section 5 and ought not to be ignored. Spatial grid: At each synchronisation cycle (Step 2 in Algorithm 1), the numerical solution of the equations (1–4) is more expensive than in a purely data-driven model, and this is not benchmarked against the retraining frequency (Section 5.6), except for a coarser mechanistic grid, where only the ML correction term (Eq. 8) is applied. Real-time inference: The trained ML layer, which involves performing the forward pass to make inferences, can be executed within the edge-computing component of Table 5 in a low-latency manner on a frequent (e.g., hourly or daily) schedule, while Eq. 6 can be naturally performed on a cloud server or on-premises server, which would require a slower cadence.

Scalability: when working at portfolio scale (as for the FHWA LTBP programme referenced in Table 4) the framework as specified is per-structure, with the possibility to derive a shared ML model with structure-specific mechanistic parameters, or to retrain a set of structure-specific models using transfer learning from a set of structure-level models as part of a base model at portfolio scale, which is consistent with the leave-one-structure-out validation already proposed in Section 5.6 but not tested at a portfolio scale. Transfer learning and federated learning: Section 2.11 points to a recurring limitation in literature – the lack of standardised, large-scale training data – and federated learning (where updates to models, not model data, are shared across structures or owners) is a literature consistent approach to this that does not require centralised data sharing agreements, particularly with regard to the durability domain.

The physics-informed loss term (Eq. 8) is tailored towards reducing the amount of data that the ML layer requires relative to a purely data-driven model, consistent with the general PINN literature, and when a sensor channel is missing or flagged as an outlier (Section 5.6 preprocessing), the mechanistic-core prior (Section 5.2) can be used in the absence of that observation in the Bayesian update (Eq. 6) by simply removing that row of the observation matrix H , providing a principled solution rather than ad hoc data imputation. Here the four questions are marked as required future work (Section 10) but not solved by this conceptual work.

10. Conclusion and Future Work

This paper has compiled evidence from 73 sources, spanning degradation mechanisms, microstructural characterisation, materials to improve durability, classical and probabilistic service-life models, and machine-learning/digital-twin structural health monitoring, to highlight a specific and literature-backed research gap: the lack of an explainable, sensor-integrated digital-twin framework that couples multi-hazard degradation models with field-updatable machine learning for cementitious materials. This revision has successfully tested that gap against six existing digital-twin and ML-durability frameworks (Table 3a), further developed the governing mathematical formulation of the mechanistic core, Bayesian updating and physics-informed ML layer (Section 5), specified a complete synchronisation algorithm and machine-learning workflow (Sections 5.1 and 5.6), extended the explainability methodology to a stated technique-selection rationale and computation procedure (Section 5.5), and introduced an explicit treatment of uncertainty propagation (Section 5.7) and a discussion of practical implementation (Section 8). To address this gap, a five-layer conceptual framework has been presented, with literature-derived parameter ranges presented with an explicit, rubric-based scoring methodology (Section 6.1); and a staged validation approach has been proposed that involves the use of named, existing laboratory and field data sources, such as the FHWA LTBP programme and the Itaipu Dam core record (Section 7).

Future work should include (i) building and sharing a standardised dataset of accelerated-exposure and field-monitoring data, directly addressing the finding of RILEM TC 315-DCS [5] regarding data-scarcity, and the federated-learning approach suggested

in Section 9.1; (ii) testing the physics-informed ML layer outlined in Section 5.4 against the qualitative comparison in Fig. 3, using real datasets of multi-mechanism data, rather than synthetic; and (iii) pilot field deployment on an instrumented marine or de-icing-salt-exposed structure to validate the Bayesian updating scheme in Section 5.3 and the practical implementation considerations outlined in Section 8, including sensor drift correction and BIM integration. Genuinely field-validated, interpretable, multi-hazard service-life prediction can only be achieved through such transparently, mathematically validated development, building on the mathematical and algorithmic specifications provided herein.

Declarations

Conflicts of Interest/Competing Interests:

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

Research involving Human participants and/or animals

Not applicable.

Ethics, Consent to Participate, and Consent to Publish declarations

Not applicable.

Clinical Trial Number

Not applicable,

Funding:

Not Applicable

Author Contribution

D.O. conceived the study, performed the literature review, developed the conceptual framework, formulated the mathematical models and computational workflow, prepared the figures and tables, and wrote the original manuscript draft. P.S. contributed to the development of the machine learning, explainable AI, and digital twin methodologies, critically reviewed the technical content, validated the computational framework, and revised the manuscript for important intellectual content. Both authors reviewed, edited, and approved the final version of the manuscript.

Acknowledgement

The authors gratefully acknowledge the support and facilities provided by the Director, Amity School of Engineering and Technology (ASET), Amity University Uttar Pradesh, Lucknow Campus, India, for facilitating this research.

Data availability statement:

Not applicable

References

1. Review of research on prediction models for residual life of concrete structures. *Coatings*. 2025;15(6):693. <https://doi.org/10.3390/coatings15060693>.
2. An integrated service. life prediction model for chloride-induced corrosion in reinforced concrete structures. *Eur J Environ Civil Eng*. 2025;30(1):1–22. <https://doi.org/10.1080/19648189.2025.2551982>.
3. Rajender A, Samanta AK, Paral A. Comparative study of corrosion-based service life prediction of reinforced concrete structures using traditional and machine learning approach. *Int J Struct Integr* 2025, 16(3), 591–621. <https://doi.org/10.1108/IJSI-02-2024-0018>
4. Characterization of mechanical property. evolution and durability life prediction of engineered cementitious composites under frozen state. *Materials*. 2025;18(10):2375. <https://doi.org/10.3390/ma18102375>.
5. Machine learning in. concrete durability: challenges and pathways identified by RILEM TC 315-DCS towards enhanced predictive models. *Mater Struct*. 2025. <https://doi.org/10.1617/s11527-025-02664-3>.
6. Finite element analysis of. chloride ion penetration and service life prediction in concrete with supplementary cementitious materials. *Sci Rep*. 2025. <https://doi.org/10.1038/s41598-025-15925-6>.
7. Sheng R, Wang K, Wei H, Lu H, Li C. Study on the performance and service life prediction of corrosion-resistant concrete cut-corner square piles. *Materials*. 2025;18(16):3776. <https://doi.org/10.3390/ma18163776>.
8. Mechanistic modeling for. coupled chloride-sulfate attack in cement-based materials. *Constr Build Mater* 2024. <https://www.sciencedirect.com/science/article/abs/pii/S0950061824043733>
9. Hydro-chemo-mechanical. coupled modeling of sulfate attack in unsaturated concrete under moisture cycling and loading effects. *Constr Build Mater* 2025. <https://www.sciencedirect.com/science/article/abs/pii/S0950061825021208>
10. Degradation of concrete in. marine environment under coupled chloride and sulfate attack: a numerical and experimental study. <https://www.researchgate.net/publication/361227002>
11. Modeling of concrete deterioration. under external sulfate attack and drying-wetting cycles: a review. <https://pmc.ncbi.nlm.nih.gov/articles/PMC11243601/>
12. Creep. chloride, carbonation and sulphate attack on concrete. <https://www.researchgate.net/publication/359143818>
13. Zhang Y, Qin Y, Guo Z, Li D. Research on performance deterioration of multi-walled carbon nanotube-lithium slag concrete under the coupling effect of sulfate attack and dry-wet cycles. *Materials*. 2023;16(14):5130. <https://doi.org/10.3390/ma16145130>.
14. Understanding the degradation. mechanisms of cement-based systems in combined chloride-sulfate attack. *Constr Build Mater* 2022. <https://www.sciencedirect.com/science/article/abs/pii/S000888462200357X>
15. A study on the application of artificial intelligence in. digital twin monitoring of building structural health. *Proc. 2025 Int. Conf. on AI and Smart Manufacturing*. <https://doi.org/10.1145/3756423.3756560>
16. A fuzzy-. set-based joint distribution adaptation method for regression and its application to online damage quantification for structural digital twin. *arXiv:2211.02656*. <https://arxiv.org/pdf/2211.02656>
17. A generative adversarial. network optimization method for damage detection and digital twinning by deep AI fault learning: Z24 Bridge structural health monitoring benchmark validation. *arXiv:2511.00099*. <https://arxiv.org/pdf/2511.00099>
18. Real-time structural. health monitoring with Bayesian neural networks: distinguishing aleatoric and epistemic uncertainty for digital twin frameworks. *arXiv:2512.03115*. <https://arxiv.org/pdf/2512.03115>
19. Mengesha G. Digital twin-driven structural health monitoring: emerging paradigms, simulation strategies and predictive intelligence. *SSRN Electron J* 2025. <https://www.researchgate.net/publication/403024735>
20. Machine learning (ML). and deep learning (DL) in sustainable concrete construction: review, trend and gap analyses. *J Asian Archit Building Eng*. 2025. <https://doi.org/10.1080/13467581.2025.2574571>.
21. Yingqun Z, Zhuang L, Nana Z. Artificial intelligence-driven health monitoring and early warning technology for concrete structures. *Int J High Speed Electron Syst*. 2026;35(03):2540450. <https://doi.org/10.1142/S0129156425404504>.

22. Sarma IV, Chanda S, Reddy MS. The convergence of artificial intelligence and structural health monitoring: a comprehensive review of methodologies, advancements, and future trajectories. *Life Cycle Reliab Saf Eng*. 2025;14(4):531–43. <https://doi.org/10.1007/s41872-025-00364-z>.
23. Li F, Luo D, Niu D. Data-intelligence driven methods for durability, damage diagnosis and performance prediction of concrete structures. *Commun Eng*. 2025;4(1):100. <https://doi.org/10.1038/s44172-025-00431-4>.
24. Flexural performance and microstructural characterization of microbially enhanced cement-reduced mortars. *Buildings*. 2026;16(5):1045. <https://www.mdpi.com/2075-5309/16/5/1045>.
25. Microstructural characterization of 3D printed cementitious materials. <https://pmc.ncbi.nlm.nih.gov/articles/PMC6766306/>
26. Georget F, et al. A critical catalogue of SEM-EDS multispectral maps analysis methods and their application to hydrated cementitious materials. *J Microsc*. 2024. <https://doi.org/10.1111/jmi.13245>.
27. edxia. microstructure characterisation from quantified SEM-EDS hypermaps. *Cement and Concrete Research* 2021. <https://www.sciencedirect.com/science/article/pii/S0008884620316070>
28. Cement-based grouting materials modified with GO/NS hybrids. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC12608761/>
29. Performance evolution. and degradation mechanism of chemically bonded phosphate ceramic cement under freeze-thaw cycles. *Materials*. 2025;18(23):5298. <https://doi.org/10.3390/ma18235298>.
30. A review of micro-. resolved crystochemical and mechanical probes for sustainable cement-based material studies. *npj Mater Sustain*. 2025. <https://doi.org/10.1038/s44296-024-00044-6>.
31. Three-dimensional visualization. of X-ray micro-CT with large-scale datasets: efficiency and accuracy for real-time interaction. *arXiv preprint*. <https://arxiv.org/pdf/2601.15098>
32. Synergistic effects of SAE/CaCO₃. on enhancing sulfate resistance in high-strength cementitious composites. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC12491487/>
33. Additive manufacturing and characterization. of architected cement-based materials via X-ray micro-computed tomography. *arXiv:1808.00396*. <https://arxiv.org/pdf/1808.00396>
34. Effects of recycled and supplemented cementitious materials on corrosion resistance. and mechanical properties in reinforced concrete. *J Compos Sci* 2025, 9(9), 457. <https://doi.org/10.3390/jcs9090457>
35. Effect of supplementary cementitious materials on corrosion resistance and durability. *Civil Environ Eng*. 2025. <https://doi.org/10.2478/cee-2025-0080>.
36. Effect of silica fume. and fly ash admixtures on the corrosion behavior of AISI 304 embedded in concrete exposed in 3.5% NaCl solution. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6926658/>
37. Mechanical performance. and corrosion resistance of concrete with partial fly ash substitution. *J Building Pathol Rehabilitation*. 2025. <https://doi.org/10.1007/s41024-025-00765-8>.
38. Dang VQ, Chau VN, Sang NT, Thuc PH, Ho LS. Mechanical properties and durability of concrete containing fly ash and GGBS for marine environment: a comprehensive study from laboratory perspective. 2026. <https://doi.org/10.1177/14644207251317612>
39. Review of the effects. of supplementary cementitious materials and chemical additives on the physical, mechanical and durability properties of hydraulic concrete. <https://pmc.ncbi.nlm.nih.gov/articles/PMC8672277/>
40. An innovative approach. to enhancing the strength and durability of recycled aggregate concrete through fly ash-silica fume coating and rice husk ash supplementation. *Sci Rep*. 2025. <https://doi.org/10.1038/s41598-025-18138-z>.
41. Bacterial concrete. the future of self-healing and sustainable infrastructure. <https://www.sciencedirect.com/science/article/pii/S2215016125004133>
42. Javeed Y, Goh Y, Mo KH, Yap SP, Leo BF. Microbial self-healing in concrete: a comprehensive exploration of bacterial viability, implementation techniques, and mechanical properties. *J Mater Res Technol*. 2024. <https://doi.org/10.1016/j.jmrt.2024.01.261>.
43. Development. enhancement of self-healing bacterial concrete and its application as sustainable concrete in modern construction - a review. *Geomicrobiol J*. 2025;42(10):1076–100. <https://doi.org/10.1080/01490451.2025.2540888>.

44. Advances in microbial. self-healing concrete: a critical review of mechanisms, developments, and future directions. *Sci Total Environ.* 2024;947:174553. <https://doi.org/10.1016/j.scitotenv.2024.174553>.
45. Bacteria-powered self. -healing concrete: breakthroughs, challenges, and future prospects. *J Ind Microbiol Biotechnol.* 2024;52:kuae051. <https://doi.org/10.1093/jimb/kuae051>.
46. State-of-the. -art review of microcapsule self-repairing concrete: principles, applications, test methods. prospects *Polym.* 2024;16(22):3165. <https://doi.org/10.3390/polym16223165>.
47. State-of-art. review of bacteria-based self-healing concrete: biomineralization process, crack healing, and mechanical properties. <https://www.sciencedirect.com/science/article/abs/pii/S095006182300911X>
48. Bacterial self-healing. and mechanical strength enhancement in concrete: a comparative study of *Bacillus subtilis*, *Bacillus sphaericus*, and *Escherichia coli*. *Innovative Infrastructure Solutions* 2025. <https://doi.org/10.1007/s41062-025-02278-2>
49. Zeenat K, Daud K. Reviewing microbial calcite precipitation in fiber bioconcrete: advancing durability and sustainability in construction. *Adv Biotechnol Microbiol.* 2024;18(3):555989. <https://doi.org/10.19080/AIBM.2024.18.555989>.
50. Evaluation of concrete carbonation based on a fiber. Bragg grating sensor. <https://pmc.ncbi.nlm.nih.gov/articles/PMC10818904/>
51. Embedded-type. optical sensors for in situ monitoring of carbonation of cementitious materials. <https://www.researchgate.net/publication/368892219>
52. Investigation of the parameters. influencing progress of concrete carbonation depth by using artificial neural networks. *Materiales de Construcción.* <https://materconstrucc.revistas.csic.es/index.php/materconstrucc/article/view/2271>
53. Machine learning formulation for. predicting concrete carbonation depth: a sustainability analysis and optimal mixture design. <https://www.sciencedirect.com/science/article/abs/pii/S2352012425008501>
54. Potential. and challenges in long-term monitoring of concrete carbonation using embedded-type optical sensors. <https://www.researchgate.net/publication/385598553>
55. Data-driven method for enhanced corrosion assessment of reinforced concrete structures. arXiv:2007.01164. <https://arxiv.org/pdf/2007.01164>
56. An embedded-sensor. approach for concrete resistivity measurement in on-site corrosion monitoring: cell constants determination. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8038218/>
57. Prediction model for the carbonation of post-repair materials. in carbonated RC structures. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5459081/>
58. Potential. and challenges in long-term monitoring of concrete carbonation using embedded-type optical sensors. *Constr Build Mater* 2024. <https://www.sciencedirect.com/science/article/pii/S2214509524011112>
59. Alexander M et al. Service life design and modelling of concrete structures - background, developments, and implementation. <https://www.scielo.org.mx/scielo.php?pid=S2007-68352018000300224>
60. Life-365 service. life prediction model version 2.0. <https://www.researchgate.net/publication/281263490>
61. Edvardsen C. The consultant's view on service life design. RILEM Proceedings PRO 070. <https://www.rilem.net/images/publis/pro070-029.pdf>
62. fib Task Group 5.6. Model Code for Service Life Design. *fib Bulletin* 34. 2006. <http://ndl.ethernet.edu.et/bitstream/123456789/24612/1/Model%20Code%20for%20Service%20Life%20Design.pdf>
63. Synthesis of service life prediction for bridges in Texas. arXiv:2509.16208. <https://arxiv.org/pdf/2509.16208>
64. DuraCrete. Probabilistic Performance-Based Durability Design: Modelling of Degradation. Document BE95-1347/R4-5, The Netherlands, 1998.
65. Tuutti K. Corrosion of Steel in Concrete. Swedish Cement and Concrete Research Institute, CBI Research Report 4:82, Stockholm, 1982.
66. Papadakis VG, Vayenas CG, Fardis MN. Fundamental modeling and experimental investigation of concrete carbonation. *ACI Mater J.* 1991;88(4):363–73.
67. Bazant ZP. Physical model for steel corrosion in concrete sea structures – theory. *J Struct Div ASCE.* 1979;105(6):1137–53.

68. Saetta AV, Scotta RV, Vitaliani RV. Analysis of chloride diffusion into partially saturated concrete. *ACI Mater J.* 1993;90(5):441–51.
69. Tang L, Nilsson LO. Chloride binding capacity and binding isotherms of OPC pastes and mortars. *Cem Concr Res.* 1993;23(2):247–53.
70. Andrade C. Calculation of chloride diffusion coefficients in concrete from ionic migration measurements. *Cem Concr Res.* 1993;23(3):724–42.
71. Neville AM. *Properties of Concrete*. 4th ed. Longman: Harlow, UK; 1995.
72. Mehta PK, Monteiro PJM, *Concrete. Microstructure, Properties, and Materials*. 4th ed. New York: McGraw-Hill; 2014.
73. Thomas MDA, Bentz EC. Life-365 computer program for predicting the service life and life-cycle costs of reinforced concrete exposed to chlorides. Life-365 Consortium; 2002.

Figures




Figure 1

Proposed sensor-integrated explainable ML digital-twin framework.




Figure 2

Illustrative Fick's-law chloride ingress profiles at one year of exposure, computed using diffusion coefficients reported for w/b = 0.45–0.55 concretes [6].




Figure 3

Qualitative capability comparison across service-life prediction approaches, synthesised from the sources cited in Table 2.




Figure 4

Thematic overlap map: the identified gap lies at the intersection of coupled multi-hazard modelling, explainable ML, and field-validated digital-twin sensing.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [SensorIntegratedMLDigitalTwinCementitiousDurabilityPaperSUPPLEMENTARY.docx](#)