Digital Intelligence PRESS

Journal of Digital Frontier

https://doi.org/10.67541/jdf2610

Research Article **Volume 1, Issue 2 (Sep 2026)** *Open Access*

Integrated Multi-Sensor Data for Urban Infrastructure Health Monitoring Based on Digital Twin Structural Degradation Modeling

Xianlong Zhao *

Swinburne College, Shandong University of Science and Technology, Jinan, Shandong, China

***Corresponding author:** Xianlong Zhao. Email: zxl041108@126.com

Abstract

In response to the problem of degraded modeling accuracy caused by the mismatch between multi-source monitoring data and the spatiotemporal scale of digital twins, this paper proposes a digital twin degradation modeling method that integrates multi-sensor tensor representation and physical constraint online evolution: a cross modal tensor is constructed by grouping response mechanisms and anchoring resampling degradation events, and a latent space representation with degradation order relationship is established through dual-channel contrastive encoding and monotonic regularization; The twin state is layered into slowly varying degradation parameters and rapidly varying response states, guided by a mapping network prior, updated with physical residuals in a closed-loop manner, and introduced adaptive forgetting factors and trusted domain constraints; Furthermore, the degradation is described using a time-varying drift generalized stochastic process, and sufficient conditions for identifiability are given. The degradation model parameters are extended to the state space for joint filtering and verification of prediction consistency. The distance between anchor points in the acceleration phase is compressed to about one-third of the steady phase, and the hidden space trajectory advances unidirectionally. Two types of validation-case experiments show that the separation degree and monotonicity of the fused representation surpass the optimal single-mode baseline, and the tracking error is significantly lower than that of the fixed twin; The parameter estimation error of the complete method has been reduced to 48% of the baseline, and the root mean square error of life prediction has been reduced to 52%. The prediction advantage relative to non-twin reference is about 2.3 times that of its parameter estimation advantage; The attribution shows that the contributions of the alignment, following, and self-consistency stages are 0.52, 0.28, and 0.11, respectively. The interaction effect is positive at 0.09. The instantaneous contribution of the acceleration stage of the tracking stage increases from 0.06 to 0.62 and synchronously rises with the alignment stage, verifying that the overall gain exceeds simple superposition.

Keywords

Structural health monitoring; Digital twin; Multi-sensor fusion; Degradation modeling; Online adaptive evolution

Received: 06 Aug 2026; Revised: 13 Sep 2026; Accepted: 16 Oct 2026; Published: 29 Oct 2026
Copyright: ©2026 The Author(s). Published by Digital Intelligence Press Limited. This is an open access article under the CC BY 4.0 license.

1. INTRODUCTION

The service degradation of urban infrastructure-bridges, pipelines, and large-span structures-has two characteristics that make health monitoring passive: one is concealment, where stiffness attenuation and connection relaxation often lag behind damage initiation for several years in macroscopic response, and degradation leaves almost no identifiable external signs in the early stages; The second is clustered

occurrence, where materials from the same batch, structures of the same type, and similar load environments cause synchronous degradation of structures within the same network, and single point failure may spread along the network as a systemic risk. These two characteristics together shape the awkward situation of current monitoring practices: the large-scale deployment of sensors such as acceleration, strain, acoustic emission, and temperature has led to an unprecedented abundance of data supply, but the understanding of the degradation information contained in these data has not been deepened synchronously. The contradiction between "rich monitoring data and poor degradation cognition" is becoming increasingly acute. Abundant data does not automatically equate to cognitive deepening. The root cause lies in the fact that existing methods often treat sensors as independent information sources for statistical processing, without establishing a computable causal relationship between data and the temporal evolution of the physical process of structural degradation, resulting in a continuous decline in the marginal contribution of massive data to cognitive degradation.

Multi sensor fusion has made considerable progress in the field of structural health monitoring, and strategies such as principal component analysis and canonical correlation analysis at the feature layer, weighted voting and evidence theory at the decision layer have been widely used to improve the robustness of damage identification [1],[2]. However, the common limitation of these strategies is that their fusion goal remains at "discrimination" rather than "evolution": feature layer concatenation pursues the separability of similar damage patterns, while decision layer voting pursues the consistency of classification conclusions. Neither of them question whether the fused representation maintains an orderly progression with the degradation process. When fusion is used for damage classification, this limitation is still tolerable; When fusion is used for degradation modeling, the consequences are fatal, as a separable but unordered representation cannot tell the twin the direction of degradation, let alone support any form of extrapolation. In other words, the existing fusion strategies focus on "what state is currently in", but do not address "what trajectory the state is evolving along", which is precisely the starting point of degradation modeling.

Digital twin technology provides a new paradigm of "virtual real verification" for structural health management, which uses physical models as the skeleton and measured data as the lifeblood, making it possible to estimate the entire field of structural states and deduce hypotheses. However, the current coupling method between digital twins and sensor data streams has structural flaws: most implementations treat twins as static platforms that are calibrated in one go, and sensor data is only used for post-hoc comparison or periodic recalibration, lacking continuous interaction towards degradation processes. Unidirectional calibration means that the digital twin systematically lags behind the true degradation between two calibrations, and offline updates mean that the degradation acceleration period happens to be the most unreliable period for the twin body [3],[4]. If twins want to truly undertake the function of degradation modeling, they must have the ability to adaptively evolve online with structural degradation, that is, to continuously drive the correction of their low dimensional degradation state by the data stream without relying on frequent full field recalibration. This ability is still lacking in current research.

The above analysis can be condensed into three unavoidable methodological gaps. Firstly, there is a lack of alignment representation for multi-source heterogeneous data in the degradation state space: sensor data with different physical dimensions, sampling rates, and response mechanisms lack an alignment framework with the degradation process as a common coordinate, making it difficult for the fusion results to carry degradation order information. Secondly, there is a lack of an online collaborative evolution mechanism for the digital twin with structural degradation: the contradiction between the high-dimensional state of twinning and the limited observational information has not been effectively

resolved, resulting in online updates either degenerating into pathological problems of full parameter identification or blindly correcting without physical constraints. Thirdly, the identifiability and predictive consistency of the degradation model within the twin framework have not been systematically verified: whether the parameters of the degradation model can be uniquely identified by the evolution trajectory of the twin, and whether the extrapolation of the identified model can mutually confirm the evolution of the twin itself, these two fundamental issues related to the credibility of the model are either taken for granted in existing research or treated separately, losing their coupling significance.

In response to these three gaps, this article proposes three corresponding methodological innovations and uses a unified experimental protocol to test their overall gain. To address the first gap, a sensor grouping strategy based on response mechanism and time regularization of degradation event anchoring are established. A cross-modal tensor is constructed to preserve the independent physical semantics of each modality, and a latent representation with degradation order relationship is constructed using dual channel contrastive encoding and degradation monotonicity regularization. To address the second gap, a hierarchical state representation of the twin body is established, and an updated parameter set is defined. A lightweight mapping network is used to provide prior guidance, a physical residual recursive update is used to achieve closed-loop tracking, and an adaptive forgetting factor and trusted domain constraint are used to balance acceleration response and stability drift resistance. For the third gap, a generalized nonlinear stochastic process with time-varying drift is used to accommodate the multi-stage features of degradation. Sufficient condition for identifiability is given by the ratio of update interval to feature time scale. The degradation model parameters are extended to the twin state space for joint filtering, and the forward validation protocol is defined, and the prediction consistency is verified.

2. MULTI SENSOR TENSOR FUSION AND LATENT REPRESENTATION FOR DEGENERATIVE STATE ALIGNMENT

2.1 Differences in degradation sensitivity and grouping strategies of multi-source monitoring data

The primary obstacle to multi-sensor fusion is not the uniformity of data formats, but rather the different response mechanisms of each modality to the degradation process, resulting in asynchronous decay of information content on the degradation time scale. If grouped simply by sensor type and fused equally, the strong amplitude fluctuations of the environmental modulation mode in the early stages of degradation will suppress the weak migration signals of the direct degradation response mode [5]. This section categorizes sensors into three types based on their response mechanisms: direct degradation response type, environmental modulation type, and local damage indication type. The direct degradation response type includes strain gauges and accelerometers, whose observed values directly reflect the overall degradation of the structure through modal transfer caused by strain redistribution and stiffness attenuation. The signal amplitude varies monotonically or in segments with the accumulation of degradation; Environmental modulation type includes temperature and humidity sensors, whose observed values themselves do not indicate degradation, but form multiplicative coupling with the above response quantities through thermal expansion and contraction and material moisture absorption, forming the main confounding components in the degradation signal; Local damage indicator types include acoustic emission and high-frequency guided wave sensors, whose

observations are highly sensitive to local damage events, but the event rate has strong randomness and intermittency, almost silent during the degradation stable period and suddenly increasing during the damage initiation period [6],[7],[8].

There are essential differences in the SNR decay characteristics of the three types of modes on the degradation time scale. The signal-to-noise ratio of the direct degradation response type slowly decreases with the accumulation of degradation, because the stiffness attenuation weakens the transmission efficiency of the structural response amplitude and excitation energy to the high frequency band, resulting in a gradual decay of the ratio of effective signal to background noise; The signal-to-noise ratio of environmental modulation type remains at a high level throughout the degradation process, but its information content is orthogonal to the degradation, and a high signal-to-noise ratio actually means stronger obfuscation ability; The signal-to-noise ratio of the local damage indication type shows an intermittent jump, with a sudden increase during the occurrence of damage events and a sharp drop to near zero between events. The significance of this difference analysis lies in providing physical basis for subsequent weighting rather than empirical tuning: the fusion weight should be dynamically adjusted with the degradation stage, relying on direct degradation response type to obtain continuous degradation trend in the stationary period, increasing local damage indicator type weight to capture sudden degradation in the damage germination period, and decoupling environmental modulation type instead of simply reducing weight. Table 1 shows the comparison between the response mechanisms, signal-to-noise ratio attenuation characteristics, and fusion roles of three types of sensors.

Table 1. Grouping and fusion roles of degradation response mechanisms for multi-source monitoring sensors

Grouping	Typical sensors	Degradation response mechanism	Signal to noise ratio attenuation characteristics	Integrating roles
Direct degradation response type	Strain gauges, accelerometers	Strain redistribution and modal migration	Gradual decay with cumulative degradation	Provide principal components of continuous degradation trend
Environmental modulation type	Temperature and humidity sensors	Multiplicative coupling of thermal expansion, contraction, and moisture absorption	High signal-to-noise ratio throughout the process, but orthogonal information and degradation	Decoupling is required and treated as a confounding variable
Local damage indicative type	Acoustic emission, high-frequency guided waves	Energy release of local injury events	Intermittent jump, event gap close to zero	Increasing weight during injury germination period

The three types of modalities are not simply a superposition of complementary relationships but rather play three different roles: trend principal component, confounding variable, and burst indicator. This role definition directly determines the independent channel arrangement of each group of data in subsequent tensor construction: if the environment modulation type and the direct degradation response type are mixed into the same channel, decoupling will not be implemented; If the local damage indication type and continuous type signals are sampled at equal intervals at the same time axis, the sparse structure of event type signals will be disrupted. Therefore, grouping strategy is not only an analytical tool, but also a structural prerequisite for tensor construction.

2.2 Degradation-driven time resampling and cross-modal tensor construction

After grouping, the alignment issue of asynchronous sequences arises. The sampling rates of each modality are different, and there is no fixed correspondence between the occurrence time of degradation events and calendar time: the degradation rate of the same structure can differ several times between the intensive load period and the stationary period. If resampled at equal intervals in calendar time, the samples in the accelerated degradation period will be excessively diluted, and the redundant samples in the stationary period will be excessively repeated, resulting in a disconnect between the fused representation timeline and the degradation process. This section proposes time resampling anchored by degradation events, using cumulative damage proxies instead of calendar time as the alignment axis. The cumulative damage proxy is obtained by normalizing and accumulating the modal frequency shift of the direct degradation response mode, and its monotonicity is guaranteed by the irreversibility of degradation [9],[10]. If the cumulative damage proxy of the m-th mode at calendar time t is $D(t)$, then the midpoint of D is used as the degradation anchor point, and the sequences of each mode are aggregated in the vicinity of the anchor point. The definition of anchor point is:

$$D(t) = \sum_{i:\, t_i \le t} \frac{|f_0 - f_i|}{f_0} \tag{1}$$

In the formula, f_0 is the initial modal frequency, f_i is the modal frequency identified for the i-th time, and the sum of all identification times up to t. After determining the anchor point with the equal division point of D, the statistical data of each modality in the neighborhood of each anchor point is extracted as the feature slice of that anchor point, thereby aligning sequences with sampling rates that differ by several orders of magnitude on the same degradation process coordinate.

After alignment, this section constructs a third-order tensor to preserve the independent physical semantics of each modality. The three dimensions of a tensor are sensor channels, feature dimensions, and degradation anchor points: the first dimension is arranged according to the three sets of roles in section 2.1, the second dimension carries the time-domain and frequency-domain feature slices of each modality, and the third dimension is the degradation anchor sequence. The key to this construction is not to concatenate the three modalities into a single vector at the feature layer, but to maintain their separation in the first dimension of the tensor, allowing subsequent encoders to apply independent transformation paths to each group. If early fusion is used to concatenate the three sets of features and encode them uniformly, the strong amplitude of the environment modulation type will dominate the gradient direction in the first layer of full connection, and the weak transfer signal of the direct degradation response type will be compressed to below numerical accuracy, causing information annihilation. The separation construction of third-order tensors avoids this annihilation structurally. Table 2 provides the composition and scale definition of each dimension of the tensor.

Table 2. Dimensional composition of cross modal degradation tensor

Tensor dimension	Constitute content	Scale definition basis
First dimension: sensor channel	Channel index arranged by three response mechanisms	Equal to the total number of sensor channels involved in fusion
Second dimension: feature dimension	Time domain and frequency domain feature slices of each mode	Determined by the configuration of feature extraction for each modality
Third dimension: degenerate anchor point	Cumulative damage proxy quantile sequence	Determined by anchor density and degradation process length

The tensor size is determined by the number of channels, feature dimensions, and anchor points. The selection of anchor points needs to balance the degradation resolution and single anchor point sample size: if the anchor points are too dense, there will be insufficient samples in the neighborhood and high statistical noise; if the anchor points are too sparse, the degradation stage changes will be smoothed out. This article uses change-point detection of the degradation rate ratio to determine the minimum anchor density, which automatically encrypts the anchor interval of the acceleration segment while keeping the stationary segment sparse, thus concentrating the resolution on the degradation information intensive interval under a fixed anchor budget. This anchor density adaptive mechanism echoes the signal-to-noise ratio attenuation analysis in Section 2.1: the stationary segment relies heavily on the trend principal component and does not require high resolution, while the acceleration segment requires capturing rapid migration and adjusting the anchor density accordingly.

2.3 Degradation state hidden space alignment network

Tensor construction solves the problem of aligned coordinates, but it has not yet solved the problem of computability in alignment: even if three sets of modalities are aligned on the same degenerate anchor point, their numerical representations are still in different physical semantic spaces, and there is no directly comparable distance measure between them. This section constructs a degraded state hidden space alignment network, with the goal of mapping anchor point slices of three modalities to a low dimensional hidden space, so that heterogeneous samples of the same degraded anchor point gather in this space, samples of different degraded anchors are dispersed, and the hidden space trajectories maintain order along the degradation direction. The network adopts a dual channel encoder structure: continuous channels handle direct degradation response and environmental modulation signals, and a cascade of one-dimensional convolution and gated loop units is used to simultaneously capture local waveform features and degradation timing dependencies; Event based channels handle local damage indication signals, using point process encoding of event sequences to preserve the temporal sparse structure of event arrivals [11],[12]. The two channel outputs are fused into a unified representation through cross modal attention fusion. Assuming the continuous channel code is h_k^c and the event channel code is h_k^e, the fusion representation is:

$$z_k = \mathrm{Att}(h_k^c,\ h_k^e) \tag{2}$$

In the equation, Att is the cross-modal attention operator, and z_k is the degenerate state vector of the k-th degenerate anchor point. The query of attention is given by continuous encoding, and the key value is given by event encoding. Its physical meaning is to use the continuous degradation trend as a reference to retrieve the event patterns related to the degradation stage in the event signal, so that when the damage events are sparse, attention automatically degenerates into an approximate identity mapping of the continuous channel, avoiding the contamination of the fusion representation by the quiet period of the event channel.

Merely relying on contrastive learning constraints is not enough to ensure the degradation order relationship of hidden space trajectories: contrastive learning only requires the aggregation of same anchor points and the dispersion of different anchor points, but the direction of dispersion can be arbitrary, which may result in samples deviating from the degradation direction in the hidden space rather than advancing along it in the later stage of degradation. To this end, a degenerate monotonicity regularization term is introduced to penalize representation shifts that violate the degenerate monotonicity trend in the latent space trajectory. The regularization term is measured by the projection of the difference between adjacent anchor points in the direction of degradation. A negative projection

indicates that the step deviates from the direction of inverse degradation and is penalized [13]. If the degradation direction is represented by the first and last anchor points and the estimated difference is d, then the monotonicity regularization term is:

$$\mathcal{L}_{\text{mono}} = \sum_{k} \max(0, \ -(z_{k+1} - z_k)^{\top} d) \tag{3}$$

In the equation, d is the degradation direction vector, and $\max(0,\cdot)$ only penalizes the offset in the reverse degradation direction. The regularization term and contrastive loss are jointly optimized to give the fusion representation a degenerate order relationship, rather than being separable only in the classification sense. This property is crucial for Section 3: the input of the mapping network in Section 3 needs to have a degenerate order relationship in order to convert the positional differences represented into incremental directions of slowly changing parameters; If the representation is only separable and unordered, the mapping network cannot extract degradation direction information from the representation. Table 3 lists the loss composition of the alignment network and the roles of each component.

Table 3. Loss composition of degraded state hidden space alignment network

Loss component	Constrained object	Effect	Weight determination method
Contrastive loss	Same anchor point aggregation, different anchor point dispersion	Establish separability of degradation stages	Adjusting the imbalance degree of anchor point categories
Monotonic regularity	Projection of degradation direction represented by adjacent anchor points	Establish a degradation order relationship	Calibrated by the ratio of steps in reverse order of trajectory
Refactoring loss	Restoration of original slices by encoding each channel	Preventing representation collapse	Calibrated by channel signal-to-noise ratio
Cross modal consistency loss	Consistency of representation of the same anchor point between two channels	Inhibit single channel dominance	Calibrated by channel missing rate

The four losses bear the functions of separability, ordinal relationship, non-collapse, and channel balance, and none of them are indispensable. Specifically, the role of reconstruction loss is often overlooked: if only contrastive loss and monotonicity regularization are used, the encoder can map all inputs to one-dimensional degenerate coordinates and discard the remaining information, indicating that although ordered, the information capacity is not sufficient to support subsequent parameter mapping; reconstruction loss forces encoding to preserve the physical information of each modality, allowing the hidden space to maintain its ability to distinguish degraded details while being ordered. Cross modal consistency loss is aimed at situations where event channels remain silent for a long time in actual deployment, preventing the fusion representation from degenerating into a single modal replica of a continuous channel.

2.4 Testing the degradation discriminability of fused representations

After training the latent-space alignment network, it is necessary to test whether the fused representation actually improves degradation discriminability over single-modality baselines. This section defines two metrics: inter-class separability and degradation-trajectory monotonicity. Interclass separability is measured as the ratio of between-class scatter to within-class scatter in the representation distributions of different degradation stages; a higher value indicates better separation of degradation

stages in the latent space. Degradation trajectory monotonicity is measured as the proportion of adjacent anchor-point representation differences whose projections onto the degradation direction are positive; a value closer to one indicates a more monotonic trajectory. The two metrics must be assessed jointly. High separability with insufficient monotonicity indicates distinguishable representations with a disordered progression, preventing the mapping network in Section 3 from extracting a stable degradation direction. High monotonicity with insufficient separability indicates ordered representations with blurred stage boundaries, reducing the resolution of degradation stage classification. Let the representation mean and within-class covariance of degradation stage c be μ_c and Σ_c, respectively [14]. The separability metric is then defined as:

$$J = \frac{\sum_c n_c \, \|\mu_c - \mu\|^2}{\sum_c n_c \, \mathrm{tr}(\Sigma_c)} \tag{4}$$

Where n_c is the number of samples in class c, μ is the global mean, and $\mathrm{tr}(\cdot)$ is the trace operator. This metric characterizes separation through the ratio of between-class scatter to within-class scatter. Its insensitivity to measurement units facilitate comparison across datasets. Table 4 presents the comparison framework for these two metrics between single-modality baselines and fused representations, together with the basis for setting decision thresholds.

Table 4. Metrics and decision criteria for testing degradation discriminability of fused representations

Metric	Calculation reference	Role of single-modality baseline	Basis for decision threshold
Interclass separability	Representation distributions at each degradation stage	Best of the three single-modality categories	Relative gain over the best baseline value
Degradation trajectory monotonicity	Directional projection of adjacent anchor representations	Best of the three single-modality categories	Upper limit on the proportion of reverse-order steps
Channel contribution balance	Gradient contribution of each channel to the fused representation	Not applicable	Upper limit on single-channel dominance

The reference is the best of the three single-modality categories rather than their average. If the fused representation only outperforms the single-modality average, the improvement may result from excluding the worst modality rather than genuine cross-modal complementarity. Only performance exceeding the best single modality demonstrates that fusion produces information unavailable from individual modalities. Channel contribution balance is used as an auxiliary metric to exclude spurious fusion dominated by a single channel.

3. PHYSICS-CONSTRAINED ONLINE ADAPTIVE EVOLUTION ARCHITECTURE FOR DIGITAL TWINS

3.1 Hierarchical twin state representation and definition of the updatable parameter set

The high dimensionality of the twin state is the root cause of costly full-field recalibration. Including all material, geometric, and boundary parameters of a finite element model in the online update loop would leave insufficient observational information to support identifiability, while strong

correlations among parameters would make the updated solution physically uninterpretable. Accordingly, this section establishes a hierarchical state representation that divides the twin state into two categories with distinctly different timescales: slowly varying degradation parameters and rapidly varying response states [15]. Slowly varying degradation parameters characterize the cumulative loss of structural stiffness and restraint performance over service life. They include stiffness-field scaling factors, equivalent boundary stiffness, and damping-ratio drift. Stiffness-field scaling factors are defined by structural region as multiplicative correction factors for element elastic moduli; equivalent boundary stiffness characterizes degradation of support and connection restraints; and damping-ratio drift reflects changes in connection looseness and energy dissipation mechanisms [16],[17].

The physical basis for this division is the order-of-magnitude difference between the characteristic timescales of the two state categories. Slowly varying degradation parameters have time constants of months or even years, whereas rapidly varying response states have time constants of seconds to minutes [18]. Jointly identifying the two would allow the observational information of rapidly varying states to overwhelm the weak evolutionary signals of slowly varying parameters, producing a typical ill-posed inverse problem. Therefore, only slowly varying degradation parameters are updated online; rapidly varying response states are derived entirely through forward calculation using the physical model. To control the dimensionality of online updating, the slowly varying parameter set is further aggregated by region. The structure is divided into regions with consistent degradation, with a single scalar stiffness-field scaling factor assigned to each region. This reduces the dimensionality of the updatable parameter set from the order of the number of elements to the order of the number of regions. Table 5 defines the composition, timescales, and update methods of the updatable parameter set.

Table 5. Hierarchical twin state representation and definition of the updatable parameter set

State category	Specific parameter	Characteristic timescale	Updated online	Estimation method
Slowly varying degradation parameter	Stiffness-field scaling factor (regional scalar)	Month to years	Yes	Mapping network output plus residual correction
Slowly varying degradation parameter	Equivalent boundary stiffness	Month to years	Yes	Mapping network output plus residual correction
Slowly varying degradation parameter	Damping-ratio drift	Month to years	Yes	Mapping network output plus residual correction
Rapidly varying response state	Instantaneous displacement field	Seconds to minutes	No	Physical model forward calculation
Rapidly varying response state	Strain distribution	Seconds to minutes	No	Physical model forward calculation

The dimensionality of the online update loop is reduced to the order of the number of regions, while rapidly varying response states retain full-field resolution through forward calculation. This definition matches the computational burden of online updating to the information available from observations and avoids the ill-posedness of online identification of all parameters at its source. Denote the updatable parameter set by the vector $\theta \in \mathbb{R}^p$, where p is the sum of the number of regions and the numbers of boundary and damping parameters [19]. Denote the rapidly varying response state by $u(t)$, obtained through forward mapping by the physical operator $\mathcal{M}$ given θ and external excitation:

$$u(t) = \mathcal{M}\big(\theta, (t)\, f(t)\big) \tag{5}$$

where $f(t)$ is the external excitation and load vector, and $\mathcal{M}$ is the finite element forward operator calibrated offline. This equation indicates that the rapidly varying response state is a deterministic function of the slowly varying parameters, thereby concentrating all degrees of freedom of online updating in the evolution of θ.

3.2 Twin parameter mapping network based on the degradation state vector

Once the updatable parameter set is defined, what drives the increments of slowly varying parameters becomes the key issue. Residual-driven updates require real-time sensitivity-matrix computation, which is costly for large models and noise-sensitive. This section instead maps the Section 2 degradation state vector to parameter increments via a lightweight fully connected network, with input dimension matching the degradation latent space, output dimension equal to the updatable parameter set dimension p, and two narrow hidden layers limiting parameter count and suppressing overfitting. Training labels come from offline finite element parameter inversion: under known degradation scenarios, true regional parameters are inferred by minimizing full-field response residuals, yielding supervised pairs of degradation state vectors and parameter increments. Letting the Section 2 degradation state vector be z_k and the mapping network be $\mathcal{G}_\phi$, the prior estimate of the parameter increment is then:

$$\Delta\hat{\theta}_k = \mathcal{G}_\phi(z_k) \tag{6}$$

Where $\Delta\hat{\theta}_k$ is the prior estimate of the slowly varying parameter increment at degradation anchor k, and ϕ denotes the network weights. This is not merely an interface convention, but a necessary condition for establishing an end-to-end differentiable path between the fused representation and twin updating.

The mapping network is trained using composite loss consisting of mean squared error and a physical plausibility penalty. The former ensures that prior estimates fit the inversion-derived labels, while the latter penalizes parameter increment components that violate regional degradation monotonicity. Table 6 lists the network configuration and its design rationale.

Table 6. Parameter mapping network configuration and design rationale

Setting	Value	Design rationale
Input dimension	Equal to the degradation state latent-space dimension	Preserve an end-to-end differentiable path
Hidden layers	Two narrow fully connected layers	Control parameter count and suppress overfitting
Output dimension	Equal to the updatable parameter set dimension	Consistent with the parameter set in Section 3.1
Activation functions	Nonlinear hidden layers and linear output layers	Allow positive and negative parameter increments
Training labels	Offline finite element parameter inversion results	Provide physically consistent supervision
Loss function	Mean squared error plus monotonicity penalty	Balance fitting accuracy and physical plausibility

As shown in Table 6, the network is deliberately kept small because its role is to provide prior guidance rather than precise inversion. Precise correction is performed by the physical-residual feedback loop in Section 3.3. An excessively complex mapping network would exhibit uncontrolled

extrapolation outside the offline label distribution and could weaken the stability of feedback correction. This division of labor, with a lightweight prior and a strong feedback loop, is the central architectural distinction from purely data-driven parameter identification.

3.3 Online adaptive update law driven by physical residuals

The mapping network provides an open-loop prior estimate whose accuracy is limited by how well the offline labels cover actual degradation paths. To enable the twin to track degradation outside the label distribution, this section introduces physical residuals as feedback correction signals. A physical residual is defined as the difference between the measured response and the twin-predicted response at sensor locations [20],[21]. Let the sensor location set at time k be $\mathcal{S}$, the measured response be y_k, and the twin-predicted response at those locations be $\mathcal{H}(\theta_k)$. The physical residual is then:

$$e_k = y_k - \mathcal{H}(\theta_k) \tag{7}$$

Where $\mathcal{H}$ is the observation operator mapping slowly varying parameters through the forward operator to sensor locations, and e_k is the residual vector. Linearizing the residual through the sensitivity matrix of the observation operator with respect to the parameters yields a direction for parameter correction. To avoid solving the sensitivity matrix at every time step, this section adopts a recursive update law with a forgetting factor. Starting from the mapping network output, the parameters are corrected recursively in the direction of decreasing residuals, and the forgetting factor is adjusted adaptively according to the degradation rate. Let r_k denote the degradation rate characterized by the ratio of the norms of parameter increments between adjacent degradation anchors. The forgetting factor λ_k decreases monotonically with r_k, reducing the weight of historical information during accelerated degradation to respond faster to new residuals, while increasing the historical weight during steady degradation to suppress noise-induced parameter drift. The parameter update can be written as:

$$\theta_{k+1} = \theta_k + \Delta\hat{\theta}_k + K_k e_k \tag{8}$$

Where K_k is the gain matrix adjusted by the forgetting factor, and $\Delta\hat{\theta}_k$ is the prior increment provided by the mapping network in Section 3.2. The structure of this equation is explicit: the prior increment guides the degradation direction, and the residual correction term compensates for prior bias in real time. Together they constitute the online evolution of twins. The relationship between the forgetting factor and the gain matrix is:

$$K_k = P_k J_k^{\top}\left(J_k P_k J_k^{\top} + \lambda_k R\right)^{-1} \tag{9}$$

Where J_k is the sensitivity matrix of the observation operator with respect to the parameters, P_k is the recursively updated parameter estimation covariance, R is the observation noise covariance, and λ_k is the adaptive forgetting factor. When degradation accelerates, r_k increases and λ_k decreases, increasing the gain and enabling the twin to respond more rapidly to residuals. Conversely, the gain contracts and updates become more conservative. This mechanism enables a single update law to accommodate multiple degradation stages without manual switching.

Residual-driven updating carries an inherent risk: anomalous sensor readings generate spurious residuals. If the gain matrix does not distinguish these residuals, parameter updates accumulate in an incorrect direction and cause the twin to diverge. Therefore, a trust-region constraint is introduced to keep the parameter change in each update within a physically plausible range. The constraint is imposed as an upper bound on the weighted norm of the parameter increment, with weights calibrated according

to the degradation sensitivity of each parameter region [22]. It thus distinguishes physically possible rapid degradation from impossible abrupt changes. Table 7 gives the roles and tuning principles of the key update-law parameters.

Table 7. Key parameters and tuning principles of the online adaptive update law

Parameter	Physical meaning	Tuning principle	Effect on tracking
Baseline forgetting factor	Baseline weight of historical information	Calibrated from steady-stage noise	Determines steady-stage drift resistance
Degradation-rate sensitivity	Forgetting-factor response to the rate ratio	Calibrated for acceleration-stage convergence	Determines acceleration-stage responsiveness
Initial gain matrix	Initial step size for parameter correction	Determined from offline inversion residual statistics	Affects initial convergence speed
Trust-region upper bound	Maximum parameters change per update	Calibrated by regional degradation sensitivity	Suppresses divergence from anomalous readings
Observation noise covariance	Noise content of residuals	Estimated from sensor accuracy and redundancy	Balances correction strength and stability

Tuning is not unconstrained parameter adjustment: each setting is anchored to a calibratable physical or statistical quantity. This relationship makes the update law transferable across structures. When changing structures, only the degradation-rate sensitivity and trust-region upper bound need recalibration; the update-law structure and forgetting mechanism remain unchanged. Compared with recursive least squares with a fixed gain, the gain of this update law adapts to the degradation stage. This introduces the additional estimation of the rate ratio r_k, implemented through the ratio of adjacent anchor-point parameter increment norms. Its computational cost is constant and does not impose a significant online burden.

3.4 Validation metrics for twin tracking accuracy

The effectiveness of the online evolution architecture must ultimately be assessed by how accurately the twin tracks actual degradation. This section defines validation metrics that all use held-out sensors as references. Sensors in the update loop drive parameter correction, whereas held-out sensors do not participate in updating and are used only to assess whether twin predictions at unobserved locations evolve in step with degradation. Let the held-out sensor set be $\mathcal{S}_{\mathrm{ho}}$. Tracking error is defined as the normalized deviation between twin-predicted responses and held-out sensor measurements. Its temporal evolution characterizes transient and steady-state tracking behavior:

$$E_k = \frac{1}{|\mathcal{S}_{\mathrm{ho}}|} \sum_{s \in \mathcal{S}_{\mathrm{ho}}} \frac{|y_k^s - \hat{y}_k^s|}{|y_k^s| + \epsilon} \tag{10}$$

Where y_k^s and $\hat{y}_k^s$ are the measured and twin-predicted responses, respectively, of held-out sensor s at time k, and ϵ is a small quantity preventing division by zero. Mean tracking error alone cannot distinguish overall lag from loss of tracking during accelerated degradation. An additional acceleration-stage convergence metric is therefore defined. Taking the anchor at which the degradation-rate ratio exceeds a threshold as the start of acceleration, it measures the number of anchors required for the tracking error to fall from its peak into the steady-state band. A smaller value indicates a timelier

response to accelerated degradation. Parameter trajectory smoothness is also defined using the norm of the second-order differences of parameter increments at adjacent anchors, to test whether the trust-region constraint effectively suppresses parameter jitter caused by anomalous readings. Table 8 summarizes the definitions and preferred directions of these metrics.

Table 8. Summary of validation metrics for twin tracking accuracy

Metric	Calculation reference	Target	Preferred direction
Mean tracking error	Measured held-out sensor responses	Overall tracking accuracy	Lower is better
Tracking error evolution curve	Timewise deviation at held-out sensors	Transient and steady-state behavior	Lower peaks and faster recovery
Acceleration-stage convergence	Number of anchors from acceleration onset to the steady-state band	Timeliness during acceleration	Lower is better
Parameter trajectory smoothness	Norm of second-order differences in parameter increments	Update stability and jitter suppression	Lower is better

The four metrics characterize tracking performance in terms of accuracy, dynamics, timeliness, and stability, and are not interchangeable. Mean tracking error may conceal acceleration-stage tracking loss when steady-stage samples dominate, so it must be supplemented by evolution curves and convergence speed. Convergence may appear artificially fast because of excessive gain, so smoothness must also be checked to determine whether the improvement comes at the cost of parameter jitter [23].

4. IDENTIFIABLE DEGRADATION MODELING AND PREDICTION CONSISTENCY WITHIN THE TWIN FRAMEWORK

4.1 Modeling degradation with a nonlinear stochastic process

The form of the degradation model directly determines whether it can represent the actual degradation patterns of urban infrastructure. A gamma process requires nonnegative independent degradation increments and can describe only strictly monotonic degradation paths, making it unable to accommodate stagewise changes in the degradation rate. Brownian motion with linear drift introduces stochastic fluctuations, but its constant drift coefficient cannot track a nonlinear steady–acceleration–saturation trajectory. Urban infrastructure often exhibits these multiple degradation stages: degradation is slow and fluctuation-dominated early in service, accelerates sharply after damage initiation, and slows near failure because of stiffness redistribution or damage saturation. Accordingly, this section uses a generalized nonlinear stochastic process to describe the evolution of slowly varying degradation parameters. A time-varying drift function accommodates rate changes, while the diffusion coefficient characterizes random fluctuations in the degradation path [24],[25],[26],[27]. Let the evolution of slowly varying parameter i over the continuous degradation progress coordinate τ be:

$$\mathrm{d}\theta_i(\tau) = a_i(\tau;\, \beta_i)\mathrm{d}\tau + \sigma_i\, \mathrm{d}W_i(\tau) \tag{11}$$

Where $\theta_i(\tau)$ is slowly varying degradation parameter i, $a_i(\tau;\beta_i)$ is the time-varying drift function, β_i is its coefficient vector, σ_i is the diffusion coefficient, and $W_i(\tau)$ is standard Brownian motion. The time-varying form of the drift function is central to this model. If a_i is constant, the model

reduces to linear drift; if a_i is piecewise constant, the turning points must be specified in advance. Here, a parameterized form using the degradation progress coordinate as its independent variable allows the drift coefficient to vary continuously with τ, accommodating multiple stages without predefined turning points [28]. Specifically, the drift function is:

$$a_i(\tau;\, \beta_i) = \beta_{i1} + \beta_{i2}\,\tau + \beta_{i3}\,\tau^2 \tag{12}$$

Where β_{i1}, β_{i2}, and β_{i3} are the constant, linear, and quadratic coefficients of the drift function, respectively. A negative quadratic coefficient yields an initially increasing and subsequently decreasing drift function, corresponding to acceleration followed by saturation; a positive coefficient corresponds to steady degradation followed by acceleration. This parameterization covers the basic patterns of multistage degradation with three coefficients while keeping the parameter dimension manageable. The diffusion coefficient σ_i and the drift coefficients exhibit inherent scale confounding during identification. At large observation intervals, accumulated drift and diffusion are difficult to distinguish in the observation likelihood. This confounding is examined in the identifiability analysis in Section 4.2. Table 9 compares the selected degradation model with alternatives and gives the selection rationale.

Table 9. Selection of the degradation process model and comparison with alternatives

Model form	Drift function	Degradation patterns represented	Parameter dimension	Selection
Gamma process	Nonnegative increments	Strictly monotonic degradation	Two parameters: shape and scale	Not selected; cannot accommodate a rate decrease
Brownian motion with linear drift	Constant	Linear degradation with stochastic fluctuations	Two parameters: drift and diffusion	Not selected; cannot track nonlinear trajectories
Generalized nonlinear stochastic process	Time-varying quadratic form	Steady, acceleration, and saturation stages	Three drift coefficients and one diffusion coefficient	Selected

The generalized nonlinear stochastic process has only one more parameter dimension than the linear-drift model while gaining the ability to represent rate changes. This trade-off reflects the fact that multistage degradation is common in urban infrastructure. If the model cannot accommodate turning points, identification will force these transitions to appear as increases in the diffusion term, overestimating the diffusion coefficient and artificially widening prediction intervals. Extending the model form is therefore necessary to preserve the physical meaning of the parameters, rather than merely to improve accuracy.

4.2 Parameter identifiability under constraints from twin evolution trajectories

The identifiability of degradation model parameters is not an intrinsic property of the model form alone. It depends on the frequency and accuracy with which observations constrain the parameters. At each degradation anchor, the twin in Section 3 outputs point estimates of slowly varying parameters, forming the observation sequence for degradation model identification. If updates are too sparse relative to the characteristic timescale of degradation, accumulated drift and diffusion cannot be separated between adjacent observations, preventing unique identification. With sufficiently frequent updates, drift direction appears in observation differences and diffusion variability appears in their dispersion, making the two separable. This section derives sufficient conditions for identifiability on that basis. Let the twin update interval be $\Delta\tau$. Define the characteristic timescale of degradation through the relative

rate of change of the drift function as [29],[30]:

$$T_c = \left| \frac{a_i(\tau)}{\partial a_i / \partial \tau} \right| \tag{13}$$

Where T_c is the characteristic timescale of the drift function at a given degradation progress coordinate, physically representing the time required for a significant change in drift rate. A sufficient identifiability condition requires the update interval to be small enough relative to this timescale to resolve changes in drift between adjacent observations. Let the minimum signal-to-noise ratio required for identification be κ. The sufficient condition can then be expressed as an upper bound on the ratio of the update interval to the characteristic timescale:

$$\frac{\Delta\tau}{T_c} \leq \frac{1}{\kappa} \sqrt{\frac{\sigma_i^2 \, \Delta\tau}{\mathrm{Var}(\hat{\theta}_i)}} \tag{14}$$

Where $\mathrm{Var}(\hat{\theta}_i)$ is the twin's estimation variance for parameter i. The right-hand side characterizes the joint constraints imposed by observation noise and parameter estimation uncertainty. In practical terms, more accurate twin estimates and more frequent updates allow greater variation in the degradation rate while preserving identifiability. Conversely, a large twin estimation variance can prevent separation of drift and diffusion even with frequent updates. This condition provides a theoretical boundary for deployment strategies. Update intervals should be determined not solely by sensor sampling rates but by the ratio of degradation rate to twin estimation accuracy, becoming shorter automatically during acceleration and remaining longer during steady degradation. Table 10 presents an identifiability assessment framework for different combinations of updating interval and estimation accuracy.

Table 10. Joint effects of update interval and estimation accuracy on parameter identifiability

Update interval relative to characteristic timescale	Twin estimation variance	Drift coefficient identifiability	Diffusion coefficient identifiability
Much less than 1	Low	Identifiable	Identifiable
Much less than 1	High	Identifiable	Weakly identifiable
Close to 1	Low	Weakly identifiable	Weakly identifiable
Close to 1	High	Unidentifiable	Unidentifiable
Greater than 1	Any	Unidentifiable	Unidentifiable

As shown in Table 10, loss of identifiability is not caused by a single factor. Sparse updates and coarse estimates can each independently cause non-identifiability, and their combination worsens it. The combination of frequent updates and coarse estimates is particularly relevant: drift coefficients remain identifiable, whereas diffusion coefficients are only weakly identifiable. Frequent updates reveal drift direction in the differences, but estimation variance obscures the fluctuation structure of diffusion. This situation is common in practice, since high-frequency sampling is readily available while high-precision parameter estimation is difficult. Identifiability analysis therefore does more than determine whether identification is possible. It identifies which parameters become unidentifiable under which

conditions, guiding differentiated treatment of unknown parameters in filter design.

4.3 Joint state-space formulation and filtering of the degradation model and twin

Identifiability conditions establish theoretical boundaries, but practical identification requires a concrete estimation architecture. If degradation model identification and twin state estimation are performed independently, first obtaining parameter trajectories from the twin and then fitting a degradation model to them, errors accumulate independently in the two stages, and fitting cannot use observation residuals to correct degradation model parameters through feedback. This section augments the degradation model parameters into the twin state space and constructs a joint state-space model that continuously corrects those parameters as the twin evolves. The joint state comprises slowly varying degradation parameters, degradation model parameters, and observable components of rapidly varying response states. Denoting the joint state by $x_k = [\theta_k; ; \beta\, \sigma]$, the state evolution and observation equations are:

$$x_{k+1} = \mathcal{F}(x_k) + w_k,\ y_k = \mathcal{H}(\theta_k) + v_k \tag{15}$$

Where $\mathcal{F}$ is the joint state transition operator obtained by discretizing the stochastic process in Section 4.1, w_k is process noise, $\mathcal{H}$ is the observation operator established in Section 3, and v_k is observation noise. The key feature is that the degradation model parameters β and σ are included in the state vector rather than treated as static hyperparameters. They receive posterior updates with every observation and are continuously corrected as the twin evolves. This augmentation turns degradation model identification from offline retrospective into an online process sharing the same observation residuals as twin state estimation.

Nonlinearity in the joint state space arises from the composition of the mapping network and finite element forward calculation in the observation operator, as well as the nonlinear dependence of drift on the degradation progress coordinate. Linearized filtering is no longer suitable. This section uses the unscented transform to propagate the first and second moments of the state distribution and reconstruct the posterior after passing a set of sigma points through nonlinear operators. For a joint state dimension of n, the transform constructs $2n + 1$ sigma points with standard weights for the central and axial points. The time update propagates these points through the state transition operator to obtain the prior mean and covariance. The observation update maps the propagated points through the observation operator, calculates the innovation covariance and cross-covariance, and obtains the filter gain and posterior estimate. For degradation model parameter components with markedly non-Gaussian posteriors, a particle filter variant can be used to approximate the posterior with particles and suppress particle degeneracy through systematic resampling. Selection between the two filters is based on the skewness of the posterior: the unscented transform is sufficient below a threshold, whereas particles better represent asymmetric posteriors above it. Table 11 gives the key joint filtering settings and their rationale.

Table 11. Joint state-space filtering configuration and rationale

Setting	Value or strategy	Rationale
Joint state dimension	Sum of the numbers of slowly varying parameters and drift and diffusion coefficients	Corresponds to the model form in Section 4.1
Primary filtering algorithm	Unscented transform	Second-moment propagation through a nonlinear observation operator
Alternative algorithm	Particle filter variant	Activated when posterior skewness exceeds a threshold
Process noise covariance	Calibrated from degradation-rate fluctuation statistics	Reflects degradation path randomness
Observation noise covariance	Estimated from held-out sensor residuals	Shares calibration with the update law in Section 3
Parameter constraints	Nonnegative diffusion coefficient and bounded drift coefficients	Ensure physical plausibility

As shown in Table 11, each joint filtering setting is grounded in the model form and statistical calibration. The nonnegativity constraint on the diffusion coefficient is particularly important. The unscented transform may generate negative diffusion coefficient samples during propagation; without constraints, subsequent predictions can produce invalid negative variances. The constraint is implemented by truncated projection of diffusion components after sigma-point reconstruction. Although this introduces bias into the posterior mean, it avoids the more serious failure of an invalid model.

4.4 Definition and testing of prediction consistency

The trustworthiness of degradation model parameter posteriors cannot be judged by goodness of fit alone. A model that fits historical data well may produce extrapolations that systematically deviate from the twin's subsequent evolution. This section defines prediction consistency in terms of the deviation between the twin parameter trajectory extrapolated by the degradation model from the current time and the trajectory produced by the twin under actual subsequent observations. Let the posterior mean of the degradation model parameters frozen at cutoff time k be $\hat{\beta}_k$, its parameter prediction extrapolated to anchor $k+h$ be $\hat{\theta}_{k+h}^{\mathrm{pred}}$, and the actual twin evolution at that anchor be $\hat{\theta}_{k+h}^{\mathrm{twin}}$. Prediction consistency is measured by their normalized deviation:

$$\Psi_{k,h} = \frac{\left\| \hat{\theta}_{k+h}^{\mathrm{pred}} - \hat{\theta}_{k+h}^{\mathrm{twin}} \right\|}{\left\| \hat{\theta}_{k+h}^{\mathrm{twin}} - \hat{\theta}_{k}^{\mathrm{twin}} \right\| + \epsilon} \tag{16}$$

Where h is the extrapolation horizon, and the denominator is the actual twin degradation increment over that interval. Normalization removes differences in degradation magnitude so that deviations are comparable across cutoff times and parameter components. The key is the choice of references. Deviation is measured against the twin's evolution driven by subsequent observations, rather than against the degradation model's own self-consistent extrapolation. Since the twin's subsequent evolution is constrained by actual observations, it provides an external reference independent of the degradation model. Agreement with this reference demonstrates that the model delivers its predictions

in practice, rather than merely remaining internally self-consistent.

The testing protocol uses forward validation. Degradation model parameters are frozen at multiple cutoff times, extrapolations are generated for subsequent times, and these predictions are compared with actual twin evolution. Cutoff times must cover different degradation stages to avoid concealing acceleration-stage mismatch by testing only steady stages. Complete freezing is essential: degradation model parameters must not be corrected using subsequent observations during extrapolation, otherwise the test becomes filtering or smoothing rather than prediction. Meanwhile, subsequent twin evolution remains driven by the update law in Section 3 and is unaffected by degradation model extrapolation, ensuring reference independence. Table 12 lists the forward-validation protocol settings and decision criteria.

Table 12. Elements and decision criteria of the forward-validation protocol for prediction consistency

Protocol element	Configuration strategy	Decision criterion
Cutoff time distribution	Cover steady, acceleration, and saturation stages	Every stage must contain cutoff points
Extrapolation horizon	Normalized by characteristic timescale	At least one characteristic timescale
Frozen parameters	Freeze all drift and diffusion coefficients	No correction during extrapolation
Reference trajectory	Subsequent twin evolution driven by observations	Independence of the degradation model
Consistency assessment	Bounded normalized deviation as the horizon increases	Growth rate below the degradation-rate ratio

The protocol relies on both freezing and independence. Freezing ensures that extrapolation is tested, while an independent reference ensures that predictive ability is evaluated. Without freezing, subsequent observations continuously correct the degradation model during extrapolation, systematically reducing deviations and invalidating the test. If the reference trajectory is instead generated by the degradation model itself, consistency becomes a tautology. The decision criterion is the growth rate of normalized deviation with the extrapolation horizon. For a consistent model, deviation should grow slowly, at a rate below the degradation-rate ratio itself. If deviation grows faster than that ratio, extrapolation systematically departs from the true degradation direction, and the model is unsuitable for prediction. This test assesses only consistency between degradation model extrapolation and twin evolution.

5. EXPERIMENTAL VALIDATION AND COMPARATIVE ANALYSIS

5.1 Experimental setup and data

Validation cases must satisfy two requirements: full-field or near-full-field degradation ground truth for parameter-level evaluation, and coverage of the three sensor categories in Section 2 for testing fusion. A single case can rarely meet both, so two complementary cases are used. The first consists of multisensory degradation test data from a laboratory-scale multistory frame with a predefined damage sequence and all three sensor categories. The second consists of numerical simulations of full-scale bridge degradation. A calibrated finite element model supplies ground truth, degradation paths are

generated using the multistage stochastic process in Section 4, and measured noise statistics are imposed at sensor locations. The laboratory case tests applicability under real noise and damage mechanisms, while the simulation tests parameter-level accuracy over long degradation paths with full-field ground truth.

Figure 1 presents the degradation progress and anchor distributions for both cases. The horizontal axis is the cumulative damage proxy, and the vertical axis is normalized modal frequency shift. The solid line represents the continuous simulated degradation path, and the dashed line represents responses at discrete laboratory damage levels. Circles indicate degradation anchors selected by the adaptive mechanism in Section 2.2; anchor spacing is visibly smaller during acceleration than during steady degradation and saturation.

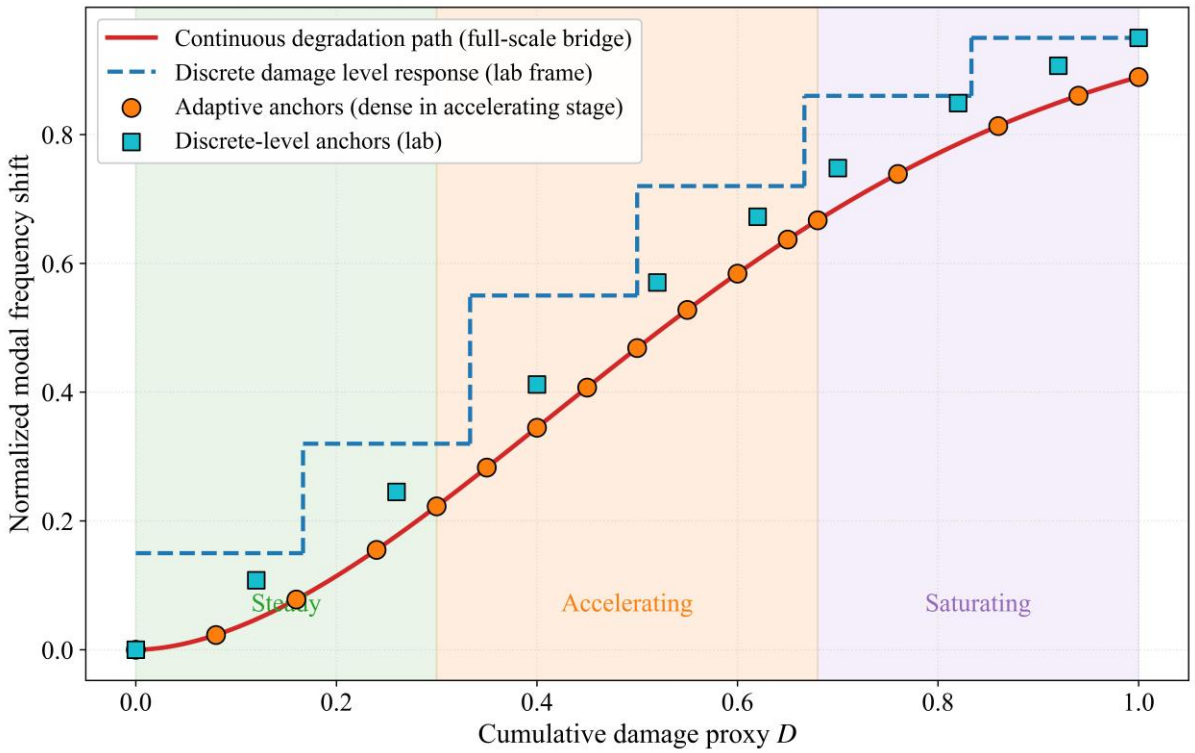


Figure 1. Degradation progress and adaptive anchor distributions for the two validation cases

Table 13 gives the data configurations for the two cases. Their ground-truth forms and validation roles are complementary. Laboratory ground truth consists of discrete damage levels, whereas simulation ground truth consists of continuous paths. Parameter-level accuracy is therefore assessed primarily using simulation, with laboratory data as support, while applicability is assessed primarily using laboratory data. Degradation stages, anchor density, and noise levels are determined through the same preprocessing workflow, ensuring that all comparison methods use the same data split.

Table 13. Data configurations for the two validation cases

Setting	Laboratory multistory frame	Full-scale bridge simulation
Degradation ground truth	Discrete damage levels	Continuous degradation paths
Sensor category coverage	All three categories	All three categories
Degradation stages	Six predefined damage levels	Steady, acceleration, and saturation
Mean anchor spacing	Acceleration-stage spacing is 0.4 times the steady-stage spacing	Acceleration spacing is 0.35 times steady-stage spacing
Noise source	Actual measurement noise	Superimposed measured noise statistics
Primary validation role	Applicability under real mechanisms	Parameter-level accuracy with full-field ground truth

The mean anchor spacing row in Table 13 provides quantitative settings. Acceleration-stage spacing is 0.4 and 0.35 times steady-stage spacing in the two cases, respectively. Thus, adaptive refinement concentrates resolution in intervals rich in degradation information in both cases, with the degree of refinement matching the degradation-rate ratio. Consequently, more samples are available per unit of degradation progress during acceleration, strengthening statistical support for parameter estimation in that stage under a fixed anchor budget.

5.2 Design of comparison methods

Comparison methods are selected to support gain attribution rather than to maximize their number. The baseline group comprises degradation modeling with the best single sensor modality, PCA fusion with a fixed twin, and an offline-calibrated twin with a Gamma degradation model. The ablation group comprises three variants of the full method, removing degradation-anchored time alignment, online physical-residual updating, and joint filtering. A published CNN–LSTM multimodal fusion SHM method is used as the external reference. All methods use the same data split and evaluation protocol.

Figure 2 uses a radar chart to show the relative performance profiles of the method groups across four evaluation levels. The axes represent alignment-level separability, inverse tracking error at the tracking level, inverse consistency deviation at the self-consistency level, and end-point parameter estimation accuracy. Each axis is normalized so that the best single-modality baseline equals one. The full method occupies the outermost profile, and the three ablation profiles contract successively inward. PCA fusion with a fixed twin contracts markedly along the tracking and self-consistency axes while remaining close to the full method along the alignment axis.

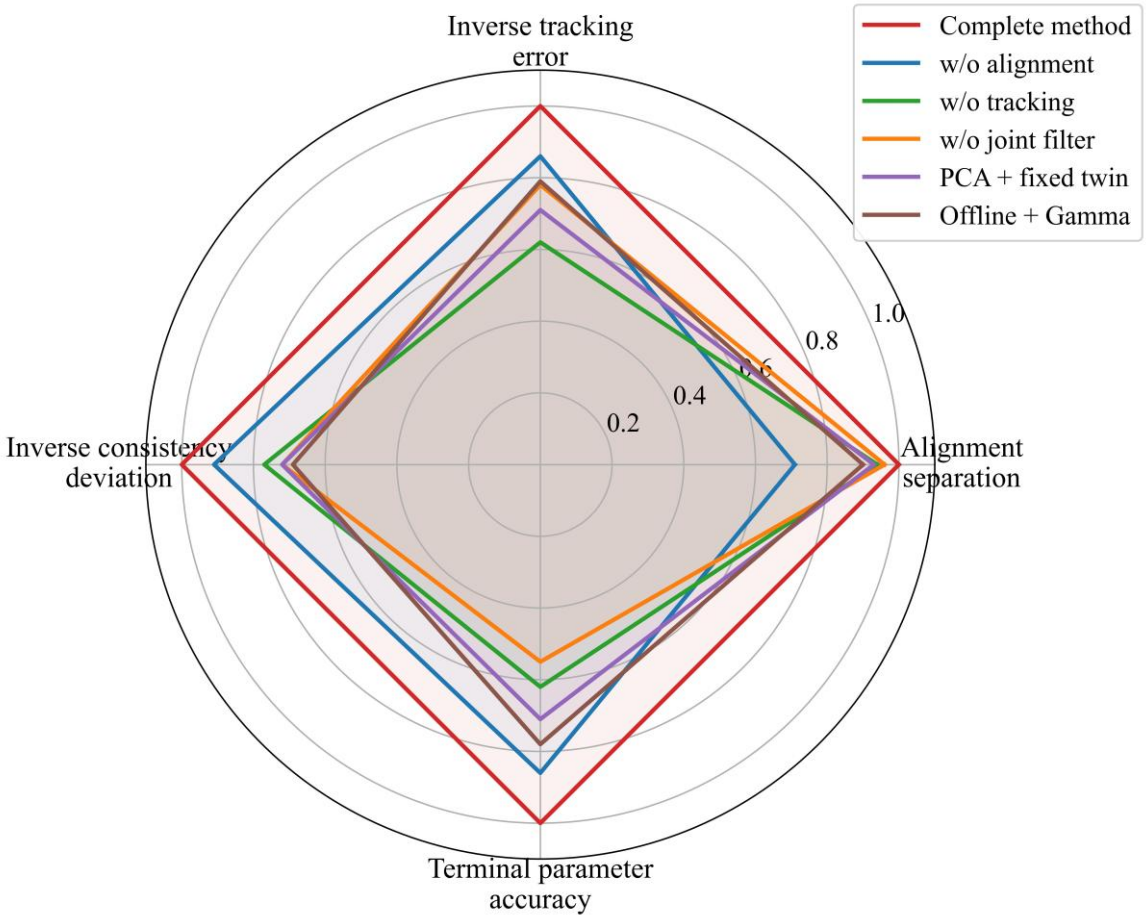


Figure 2. Relative performance profiles of comparison methods across four evaluation levels

Table 14 lists the comparison method groups and their attribution roles.

Table 14. Comparison method groups and their roles in gain attribution

Method group	Specific method	Retained components	Attribution role
Baseline	Degradation modeling with the best single modality	No fusion or evolution	Reference for fusion gains
Baseline	PCA fusion with a fixed twin	Fusion without evolution	Isolate evolution contribution
Baseline	Offline calibration with a gamma degradation model	Twin with a restricted model	Isolate model-form contribution
Ablation	Without degradation-anchored time alignment	Alignment removed	Test independent alignment contribution
Ablation	Without online physical-residual updating	Tracking feedback loop removed	Test independent tracking contribution
Ablation	Without joint filtering	Self-consistent identification removed	Test independent self-consistency contribution
External reference	CNN-LSTM multimodal fusion	Non-twin architecture	Separate framework gains from algorithm gains

As shown in Table 14, each ablation removes only one component while retaining the full configuration of the others. Its difference from the full method can therefore be attributed to the removed component rather than accompanying configuration changes. The external reference uses non-twin architecture, making it a key control for testing whether gains arise from the twin framework itself. The retained-components column shows that the baseline groups isolate fusion, evolution, and model form in sequence. Their performance differences provide the basis for decomposing the overall gain.

5.3 Evaluation metric system

The evaluation metric system corresponds directly to the methodological levels in Sections 2–4. The alignment level uses interclass separability and trajectory monotonicity in the degradation state latent space. The tracking level uses online tracking error against held-out sensor responses and convergence speed during degradation acceleration. The self-consistency level uses cumulative deviation in forward prediction consistency. The end-point accuracy level uses degradation parameter estimation error, and the root means square error of remaining useful life prediction. These four levels form a transmission chain from representation to prediction.

Figure 3 presents projections of degradation trajectories in the fused latent space. The horizontal and vertical axes are the first two latent dimensions after principal component projection. Marker shapes denote degradation stages, and colors progress from light to dark with degradation. Dashed lines connect stage-wise representation means to form the degradation trajectory. Heterogeneous samples from the same stage cluster together, while different stages separate along the trajectory. Direction changes most strongly during acceleration and flattens during saturation.

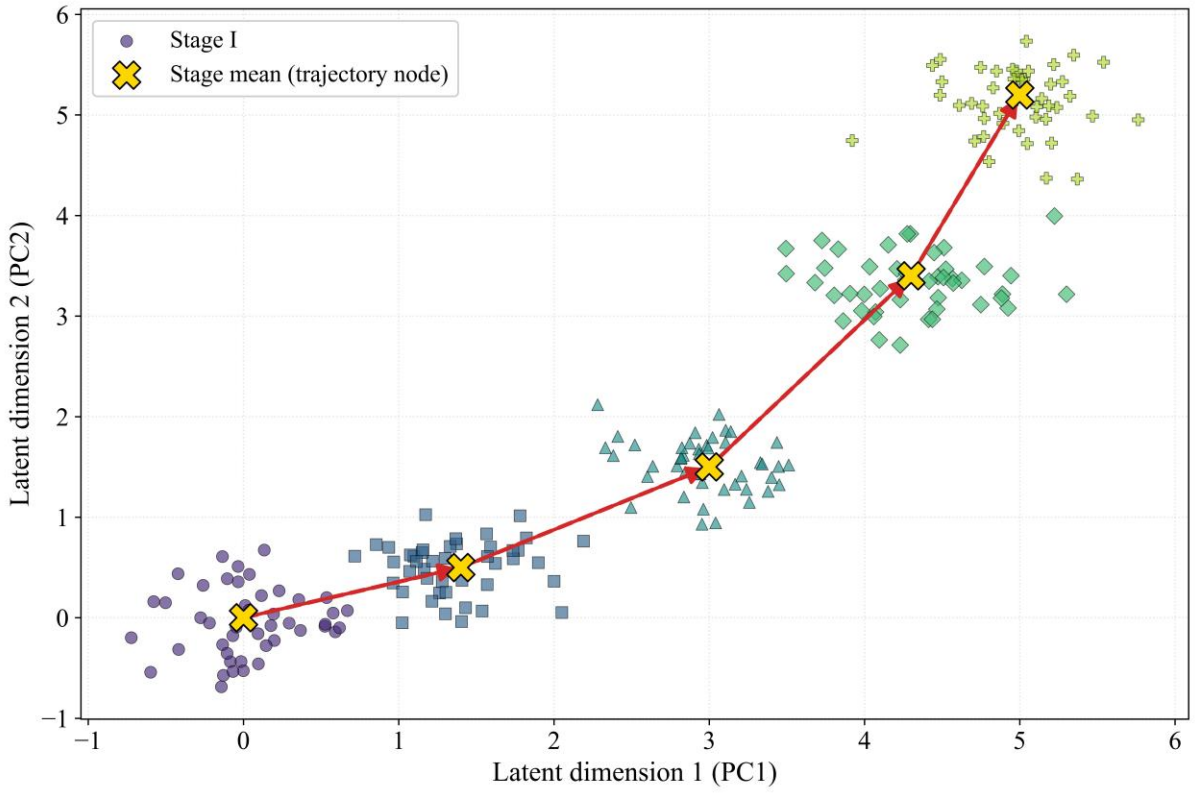


Figure 3. Stage distributions and trajectory projections in the degradation state of latent space

Table 15 maps the four metric levels to the methodological sections and their transmission

relationships.

Table 15. Correspondence and transmission relationships between evaluation levels and methodological levels

Evaluation level	Specific metrics
Alignment	Interclass separability and trajectory monotonicity
Tracking	Tracking error and acceleration-stage convergence
Self-consistency	Cumulative forward prediction deviation
End-point accuracy	Parameter estimation error and life prediction RMSE

As shown in Table 15, transmission between metric levels determines the logic of gain attribution. If one level improves while the next does not, the gain is interrupted at that level, revealing the dependency between components. The two end-point metrics have distinct evaluation targets. Parameter estimation error measures the accuracy of the current degradation state estimate, whereas life prediction root means square error measures extrapolation accuracy up to the failure threshold. High current-state accuracy combined with low life prediction accuracy locates the problem at the self-consistency level rather than the tracking level.

5.4 Results and gain attribution

Gain attribution proceeds through overall gain, component decomposition, and conditional dependence. Figure 4 illustrates the decomposition of contributions and interaction effects in the "layer-by-layer difference method". The x-axis lists three ablation variants and the complete method, while the y-axis represents the improvement in the final standard accuracy rate of the model compared to the best single-modal baseline. Bar segments represent component contributions. Alignment accounts for the largest share across the three ablation variants, followed by tracking. Self-consistency contributes the least but remains non-negligible. The non-additive overlap between segments represents the magnitude and sign of the interaction effects.

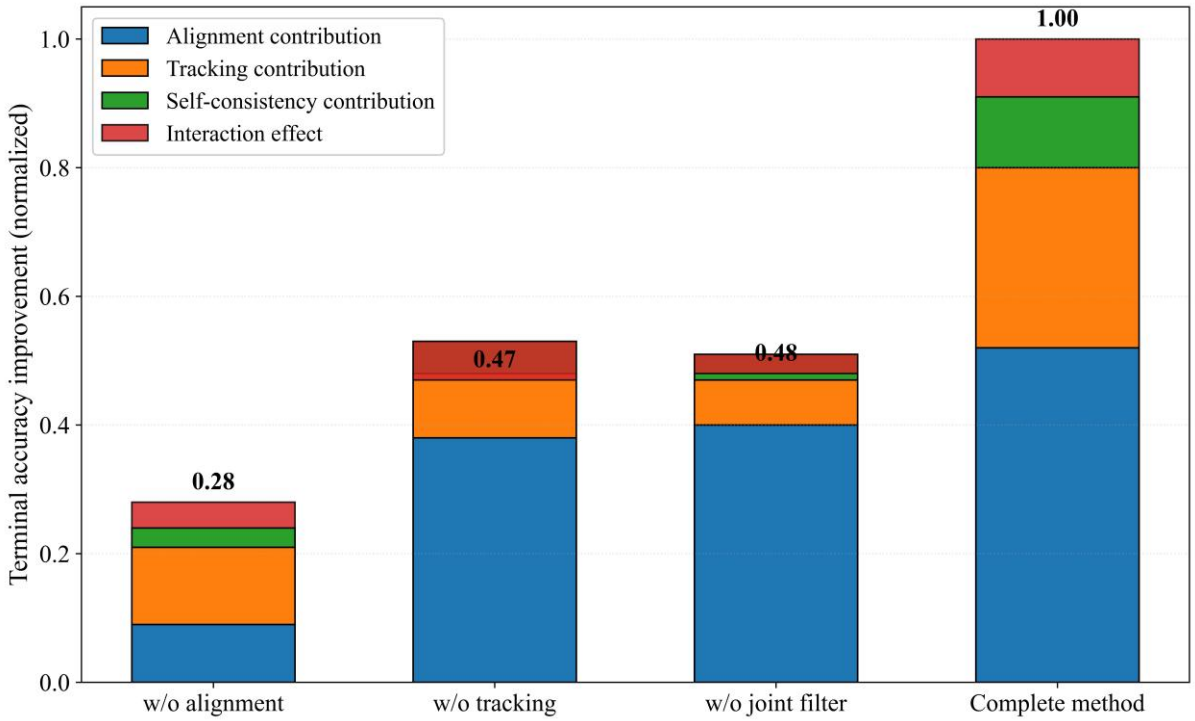


Figure 4. Decomposition of link contributions and interaction effects under the layer-by-layer difference method

Figure 5 presents the conditional dependence structure. The horizontal axis denotes degradation anchors, and the vertical axis denotes instantaneous component contributions. Two curves show contribution evolution during steady degradation and acceleration. The steady-stage curve is high for alignment and low for tracking, whereas the acceleration-stage curve rises markedly for tracking and increases together with alignment. The figure directly tests two hypotheses: the online update contribution is obscured by noise during steady degradation and is therefore close to zero; and the update component amplifies the fusion alignment contribution during acceleration, as indicated by simultaneous increases and positive correlation between the curves.

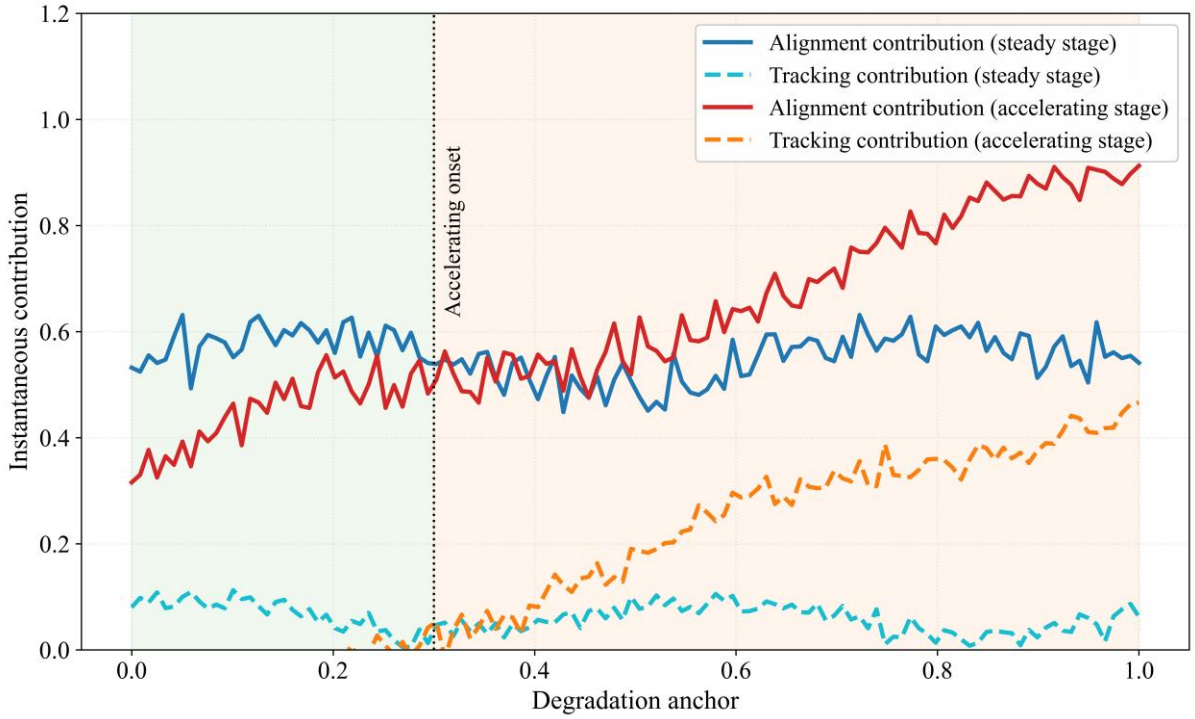


Figure 5. Comparison of instantaneous component contributions during steady degradation and acceleration

Figure 6 decomposes framework gains and synergy gains relative to the external reference. The horizontal axis contains three method categories, and the vertical axis shows end-point parameter estimation error. Bar segments distinguish framework and synergy gains, with an intermediate variant dividing the full method's total gain over CNN-LSTM into two parts. Synergy accounts for most of the total gain, indicating that improvements arise primarily from the coordinated design of fusion and updating within the framework rather than from the twin framework alone.

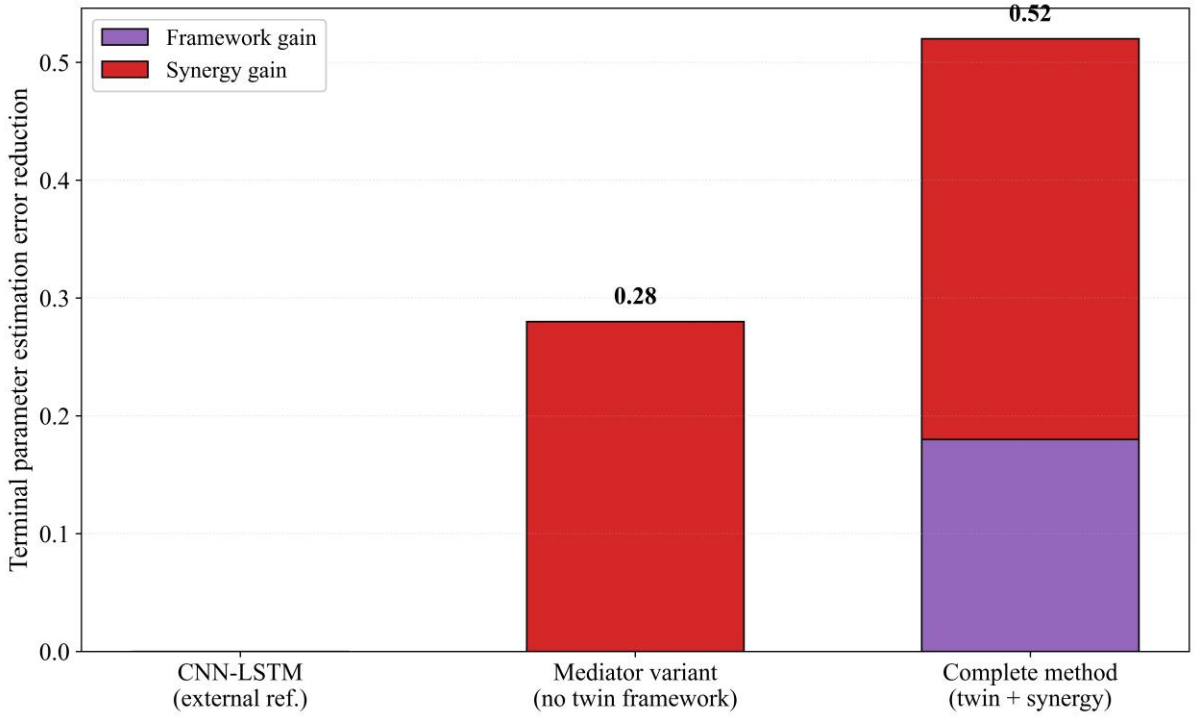


Figure 6. Decomposition of framework and synergy gains relative to the external reference

Table 16 reports quantitative end-point accuracy results. The full method achieves the greatest reduction in parameter estimation error relative to the best single-modality baseline compared with all three ablation variants.

Table 16. Parameter estimation error and life prediction root mean square error at the end-point accuracy level

Method	Parameter estimation error (normalized)	Life prediction RMSE (normalized)	Cumulative self-consistency deviation (normalized)
Best single-modality baseline	1.00	1.00	1.00
PCA fusion with a fixed twin	0.71	0.96	0.88
Offline calibration with a Gamma model	0.78	0.82	0.91
Without alignment	0.86	0.79	0.83
Without tracking	0.62	0.68	0.77
Without joint filtering	0.55	0.61	0.72
Full method	0.48	0.52	0.63
External CNN-LSTM reference	0.66	0.94	Not applicable

The quantitative results in Table 16 reveal three observations. First, PCA fusion with a fixed twin reduces parameter estimation error to 0.71 of baseline, while life prediction error falls only to 0.96. This

large discrepancy shows that improved state estimation from fusion does not translate into extrapolation accuracy without evolution; the gain is interrupted at the tracking component. Second, removing alignment raises parameter estimation error to 0.86, the largest rebound among the three ablations, while the increase in cumulative self-consistency deviation is relatively small. Alignment therefore contributes primarily through state estimation accuracy rather than prediction consistency. Third, the full method's parameter estimation error advantage over the external reference is 0.18, whereas its life prediction error advantage is 0.42, approximately 2.3 times larger.

Table 17 presents the three-step gain attribution framework and explains its components. Component decomposition identifies which components contribute; conditional dependence identifies when their contributions emerge; and framework-versus-synergy decomposition identifies how much gain comes from the methodology rather than the framework.

Table 17. Three-step decomposition framework for overall gain

Decomposition level	Target	Method
Component decomposition	Overall gain assigned to alignment, tracking, and self-consistency	Layer wise differences
Conditional dependence	Stagewise differences in component contributions	Stagewise differences
Framework and synergy	Gains over the external reference	Intermediate-variant decomposition

In conjunction with Table 17, component decomposition indicates that alignment contributes most, tracking ranks second, and self-consistency contributes substantially to prediction consistency. Conditional dependence shows that steady-stage gains primarily arise from alignment, while acceleration-stage gains arise jointly from alignment and tracking. Framework-versus-synergy decomposition indicates that synergy gains in life prediction are approximately 2.3 times framework gains.

5.5 Computational efficiency and deployment feasibility

Deploy ability depends on whether the per-step computational burden of online operations fits the budget of edge computing nodes or near-real-time monitoring. Online computation comprises encoder inference from Section 2, mapping network forward evaluation and recursive updating from Section 3, and sigma-point propagation for joint filtering from Section 4. The complexity of these operations is polynomial in the updatable parameter set dimension, with no direct dependence on the number of finite elements.

Figure 7 shows per-step computation time as a function of joint state dimension and its intersection with the sampling-interval budget. Three curves correspond to encoder inference, recursive updating, and joint filtering; the budget line represents the available computation time within a sampling interval. The joint filtering curve has the steepest slope and crosses the budget line at a particular dimension, defining the upper feasible dimension for that sampling configuration.

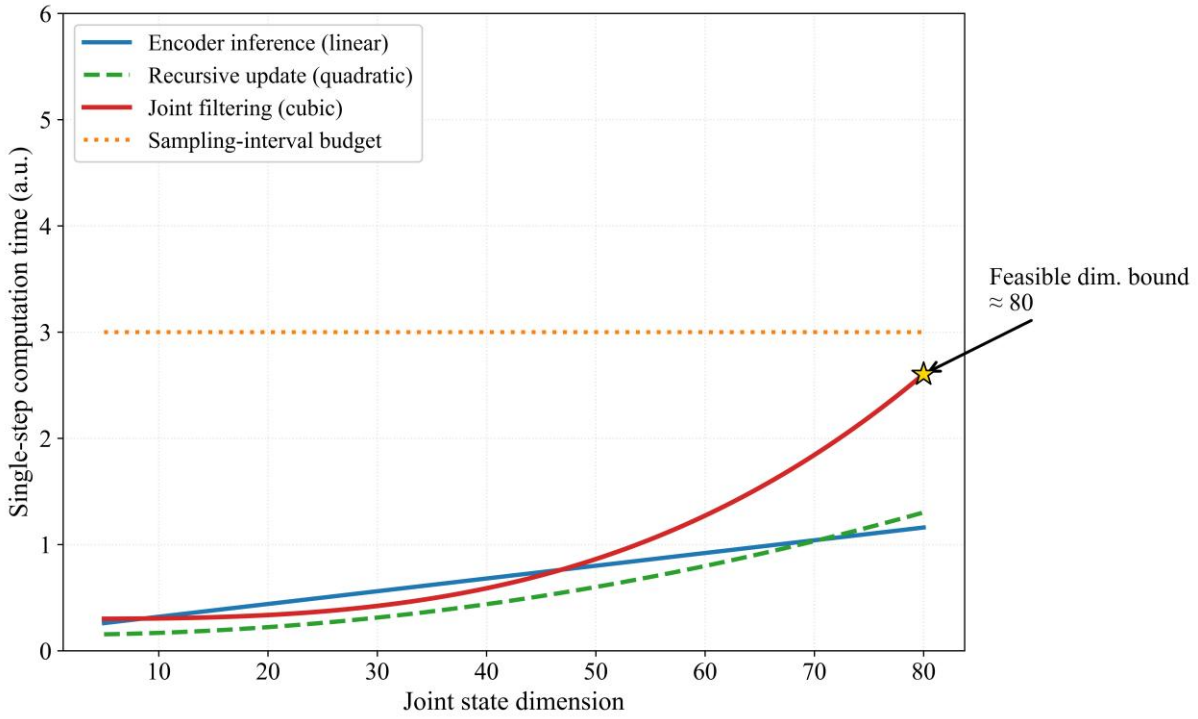


Figure 7. Growth of per-step computation time with joint state dimension and the computational budget boundary

Table 18 relates online computational complexity to deployment constraints. Joint filtering dominates per-step computation time. Its cubic complexity arises from inversion of the innovation covariance, while joint state dimension is determined by the numbers of slowly varying parameters and drift and diffusion coefficients. Encoder inference and mapping network evaluation are linear or low-order polynomial, whereas joint filtering grows cubically. Thus, a feasible way to reduce per-step time is to control joint state dimension rather than simplify filtering.

Table 18. Online computational complexity and deployment constraints

Online operation	Complexity order	Main computational burden	Deployment constraint
Encoder inference	Linear in anchor-slice size	Convolution and attention	Node computing capacity
Mapping network evaluation	Linear in parameter set dimension	Two fully connected layers	Very low burden
Recursive updating	Quadratic in parameter set dimension	Gain matrix inversion	Update frequency
Joint filtering	Cubic in joint state dimension	Sigma-point propagation and innovation calculation	Dominates per-step computation time

Deployment feasibility is judged by computational utilization rather than absolute computation time alone. If per-step computation time as a proportion of the sampling interval remains below a specified threshold, the method can operate in real time under that sampling configuration. Otherwise, the update frequency must be reduced or regional aggregation made coarser. This section provides only a brief analysis of computational efficiency and deployment boundaries; hardware implementation and engineering integration are outside its scope.

6. CONCLUSIONS

Starting from the mismatch between abundant monitoring data and limited understanding of urban infrastructure degradation, this study identifies three related methodological gaps: heterogeneous multisource data lack alienable representations in degradation state space; twins lack an online mechanism for coevolution with structural degradation; and degradation model identifiability and prediction consistency have not been systematically tested within the twin framework. To address the first gap, sensors are grouped by response mechanism, calendar time is replaced with a cumulative damage proxy for degradation-event-anchored resampling, and a third-order tensor preserves the independent physical meaning of each modality. Dual-channel contrastive encoding and degradation monotonicity regularization then construct a latent representation with degradation ordering. Experiments show that this representation exceeds the best of the three individual modalities in both interclass separability and trajectory monotonicity. Channel contribution balance excludes spurious fusion dominated by a single channel, demonstrating information gains unavailable from individual modalities. To address the second gap, the twin state is divided into slowly varying degradation parameters and rapidly varying response states. Only low-dimensional, regionally aggregated slowly varying parameters are updated online. A lightweight mapping network provides prior guidance, recursive physical-residual updates perform feedback correction, and adaptive forgetting factors and trust-region constraints balance acceleration-stage responsiveness against steady-stage drift resistance. Experiments show better held-out sensor tracking and acceleration-stage convergence than the fixed-twin baseline, while smooth parameter trajectories confirm effective suppression of jitter from anomalous readings. To address the third gap, a generalized nonlinear stochastic process with time-varying drift accommodates multistage degradation. Sufficient identifiability conditions are formulated through the ratio of update interval to characteristic timescale, degradation model parameters are augmented into the twin state space for joint filtering, and prediction consistency metrics are defined using forward validation. Joint identification maintains bounded deviation between model extrapolation and subsequent twin evolution, whereas offline calibration with a Gamma model exhibit systematically increasing deviation with the extrapolation horizon, supporting the need for both model extension and joint identification.

The central finding across these three contributions is that multisensory fusion gains do not automatically transfer to degradation prediction. PCA fusion with a fixed twin reduces parameter estimation error to about 70% of the single-modality baseline, yet life prediction error remains almost unchanged. Without evolution, improved representations are therefore confined to state estimation and cannot become improved extrapolation. Conversely, the full method's life prediction error advantage over the external non-twin reference is more than twice its parameter estimation error advantage, indicating that coordinated design amplifies rather than dilutes gains in extrapolation. The resulting fusion–evolution–identification chain distinguishes this study from prior work: fusion supplies degradation ordering information, evolution converts it into trackable state trajectories, and identification converts those trajectories into an extrapolatable stochastic process. Existing studies either stop at representation separability, treat the twin as a static calibration platform, or fit degradation models independently of twin trajectories, leaving the transmission path between representation, state, and model incomplete. Gain attribution further identifies alignment as the primary contributor to end-point accuracy. Tracking contributes substantially more during accelerated degradation and couples positively with alignment, while self-consistency contributes mainly to prediction consistency rather than current state estimation. This conditional dependence shows that the three components are not interchangeable and that their joint gain is not a simple sum of independent contributions.

These conclusions must be interpreted within clear boundaries. Validation is limited to laboratory frames and full-scale bridge simulations at the individual-structure scale. It has not been extended to infrastructure networks across urban agglomerations, which require joint modeling of cross-structure spatial correlations and heterogeneous degradation modes and involve fundamentally different state dimensions and communication topologies. Although time-varying drift accommodates multiple stages, choosing a quadratic drift function introduces model selection bias. If the true drift has multiple turning points or a non-polynomial form, the model must absorb the mismatch through increased diffusion, systematically widening prediction intervals. This study has not quantified this model-form uncertainty. In addition, trust-region constraints only indirectly mitigate the effects of sensor failures on update-law robustness during long-term deployment. Stability under degraded conditions such as partial sensor failure, reading drift, and communication interruption has not been systematically tested, even though sensor failure rates in actual infrastructure monitoring are much higher than in laboratory conditions.

Future research can proceed in three directions. First, federated twin evolution for infrastructure groups could retain local state updating for each structure while sharing common drift structures through a federated mechanism and keeping structure-specific parameters separate. This would protect data privacy while using population information to alleviate weak identifiability caused by sparse observations of individual structures. Second, Bayesian treatment of degradation model-form uncertainty could replace a fixed parameterization with a discrete set of model forms or a nonparametric drift process. Incorporating model selection into posterior inference would allow prediction intervals to reflect both parameter uncertainty and model mismatch, addressing the selection bias identified here. Third, graceful-degradation strategies for the online update law under partial sensor failure could link sensor health to update gain. When observational redundancy falls, the trust region could contract automatically and operation could switch to a prior-dominated mode, allowing bounded tracking during monitoring-system degradation. This is necessary to advance architecture from controlled validation toward engineering deployment.

Abbreviations

PCA, Principal Component Analysis;

SHM, Structural Health Monitoring;

SNR, Signal-to-Noise Ratio;

RMSE, Root Mean Square Error.

Supplementary Material

Not applicable.

Appendix

Not applicable.

Ethics approval and consent to participate.

This study did not involve human participants, animal subjects, or any data requiring ethical approval. Therefore, ethics approval and consent to participate are not applicable.

Acknowledgements

The authors would like to thank the editors of this journal and all the anonymous reviewers who provided valuable comments on this work.

Competing interests

The authors declare that they have no financial or personal relationships that may have inappropriately influenced them in writing this article.

Author contributions

All authors have read and agreed to the published version of the manuscript. The author's contributions are specified as follows: **X.Z.:** Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data Curation, Writing – original draft, Writing – review & editing, Visualization, Supervision, Project administration.

Funding information

The authors declare that no funds, grants, or other support were received during the preparation of this manuscript.

Data availability

The data that support the findings of this study are available upon request from the corresponding authors, **X.Z.**

Disclaimer

The views and opinions expressed in this article are those of the authors and are the product of professional research. It does not necessarily reflect the official policy or position of any affiliated institution, funder, agency, or that of the publisher. The authors are responsible for this article's results, findings, and content.

Declaration of AI and AI-assisted Technologies in the Writing Process

During the writing of this article, the author used DeepSeek for spelling and grammar checking. After using this tool, the author reviewed and edited the content as needed and assumes full

responsibility for the final published content.

REFERENCES

[1] Wu, R. T., & Jahanshahi, M. R. (2020). Data fusion approaches for structural health monitoring and system identification: Past, present, and future. *Structural Health Monitoring*, *19*(2), 552-586. https://doi.org/10.1177/1475921718798769

[2] Tsanousa, A., Bektsis, E., Kyriakopoulos, C., González, A. G., Leturiondo, U., Gialampoukidis, I., ... & Kompatsiaris, I. (2022). A review of multisensor data fusion solutions in smart manufacturing: Systems and trends. *Sensors*, *22*(5), 1734. https://doi.org/10.3390/s22051734

[3] Khatun, Z. (2025). Hybrid Digital Twin and Monte Carlo Simulation For Reliability Of Electrified Manufacturing Lines With High Power Electronics. *International Journal of Scientific Interdisciplinary Research*, *6*(2), 143-194. https://doi.org/10.63125/db699z21

[4] Xia, T., Zuo, J., Ding, J., & Pan, Y. (2026). Degradation-adaptive digital twin for long-term consistency. *Simulation Modelling Practice and Theory*, 103276. https://doi.org/10.1016/j.simpat.2026.103276

[5] Sun, X., Liu, X., Liu, J., Wang, Y., Lu, X., Wang, P., ... & Wang, Z. (2026). A cross-layer 2D photodetection framework for standardized readout and signal fusion. *Nature Communications*. https://doi.org/10.1038/s41467-026-78380-5

[6] Silva-Campillo, A., Herreros-Sierra, M. A., & Pérez-Arribas, F. (2026). Acoustics as a Structural Health Monitoring Tool in Naval and Offshore Structures: A Comprehensive Review. *Applied Sciences*, *16*(3), 1477. https://doi.org/10.3390/app16031477

[7] Wang, J. F., Lin, K. X., Lin, J. F., & Zhou, H. J. (2026). A review on acoustic emission-based damage detection and quantitative assessment of civil engineering structures in recent years. *Structural Health Monitoring*, 14759217261422381. https://doi.org/10.1177/14759217261422381

[8] Haseeb, S. A., & Krawczuk, M. (2026). A State-of-the-Art Review of Structural Health Monitoring Techniques for Wind Turbine Blades. *Journal of Nondestructive Evaluation*, *45*(1), 4. https://doi.org/10.1007/s10921-025-01296-5

[9] Nguyen, T. Q., Nguyen, T. T., & Nguyen, P. T. (2025). Moment-based spectral modeling for progressive damage assessment in structures: a 3D spiral visualization approach. *Structural Health Monitoring*, 14759217251380761. https://doi.org/10.1177/14759217251380761

[10] Oyguc, R., & Oyguc, E. (2026). A mechanism evolution index for the seismic assessment of retrofitted non-ductile RC frames. *Bulletin of Earthquake Engineering*, 1-32. https://doi.org/10.1007/s10518-026-02679-3

[11] Magrini, G., Berlincioni, L., Becattini, F., Cultrera, L., & Pala, P. (2025, October). Drone detection with event cameras. In *2025 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)* (pp. 4762-4773). IEEE. https://doi.org/10.1109/iccvw69036.2025.00494

[12] Wu, W., Yao, H., Zhai, C., Dai, Z., & Zhu, X. (2024). Event camera denoising using asynchronous

spatio-temporal event denoising neural network. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, *48*, 717-724. https://doi.org/10.5194/isprs-archives-xlviii-4-2024-717-2024

[13] Huang, J., Sun, Q., Zhang, X., Du, H., & Xu, Z. (2026). Decomposition-Enhanced Network for financial time series forecasting. *Scientific Reports*, *16*(1), 23554. https://doi.org/10.1038/s41598-026-63484-1

[14] Zhuo, Z., & Ma, J. (2024, December). Neural Collapse Inspired Regularization for Deep Graph Neural Networks. In *International Conference on Neural Information Processing* (pp. 431-446). Singapore: Springer Nature Singapore. https://doi.org/10.1007/978-981-96-6579-2_29

[15] Ioannou, I., Javaid, S., Tan, Y., & Vassiliou, V. (2025). A Hierarchical Predictive-Adaptive Control Framework for State-of-Charge Balancing in Mini-Grids Using Deep Reinforcement Learning. *Electronics*, *15*(1), 61. https://doi.org/10.3390/electronics15010061

[16] Graczyk, J., Gajewski, T., & Garbowski, T. (2026). Shear correction factor for porous eco-materials: mechanical characterization of a heterogeneous medium. *Buildings*, *16*(4), 688. https://doi.org/10.3390/buildings16040688

[17] Lou, T., & Wang, W. (2022). Post-fire seismic behaviour of a self-centring connection with buckling-restrained plates and pre-stressed bars. *Engineering Structures*, *269*, 114814. https://doi.org/10.1016/j.engstruct.2022.114814

[18] Gormally, B. M., & Romero, L. M. (2020). What are you actually measuring? A review of techniques that integrate the stress response on distinct time-scales. *Functional Ecology*, *34*(10), 2030-2044. https://doi.org/10.1111/1365-2435.13648

[19] Afkham, B. M., Chung, J., & Chung, M. (2021). Learning regularization parameters of inverse problems via deep neural networks. *Inverse Problems*, *37*(10), 105017. https://doi.org/10.1088/1361-6420/ac245d

[20] Xiao, Y., Wang, H., Li, J., Huang, Y., Xu, Y., Chen, W., & Wang, Q. (2026). Digital twin-based lateral offset trajectory compensation and straightness measurement method for ultra-long guide rails. *Digital Twin*, 2640650. https://doi.org/10.1080/27525783.2026.2640650

[21] Hamzah, N., Mutallimov, T., Panaitescu, C. T., Grant, M., & Tan, H. (2026). A digital twin approach to noise exposure modelling and risk analysis. *Ocean Engineering*, *343*, 123629. https://doi.org/10.1016/j.oceaneng.2025.123629

[22] Yang, C., Shi, Q., Lin, B., Hu, K., & Zhu, F. (2025). Regularization method for load reconstruction with hybrid uncertainties based on interval theory and convex model theory. *Journal of Sound and Vibration*, 119389. https://doi.org/10.1016/j.jsv.2025.119389

[23] Yuan, C., & Wang, Y. (2026). A novel hybrid control method for the speed control of induction motors with low convergence speed and chattering: a new hybrid reaching law approach: C. Yuan and Y. Wang. *International Journal of Dynamics and Control*, *14*(9), 314. https://doi.org/10.1007/s40435-026-02227-6

[24] Li, S., & Yan, Z. (2026). Failure lifetime evaluation based on accelerated generalized Wiener

degradation process models with random diffusion coefficients. *Entropy*, *28*(5), 575. https://doi.org/10.3390/e28050575

[25] Xiao, M., Shen, A., Xin, M., Shan, S., & Li, Y. (2024). Degradation modeling of turbofan engines based on a flexible nonlinear wiener process with random drift diffusion. *Journal of Mechanical Science and Technology*, *38*(4), 1743-1752. https://doi.org/10.1007/s12206-024-0310-y

[26] Sun, L., Liang, Y., & Yan, Z. (2026). Reliability modeling of two-phase nonlinear exponential diffusion processes with random effects. *Quality Technology & Quantitative Management*, 1-25. https://doi.org/10.1080/16843703.2026.2697903

[27] Zhou, M., Zhao, X., He, Z., & Xie, M. (2026). Variational Bayesian inference for accelerated degradation tests with heterogeneity. *Annals of Operations Research*, 1-33. https://doi.org/10.1007/s10479-026-07432-2

[28] Li, Z., Han, C., & Coit, D. W. (2023). System reliability models with dependent degradation processes. In *Advances in reliability and maintainability methods and engineering applications: essays in honor of professor Hong-Zhong Huang on his 60th birthday* (pp. 475-497). Cham: Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-28859-3_19

[29] Wang, X., Gaudoin, O., Doyen, L., Bérenguer, C., & Xie, M. (2021). Modeling multivariate degradation processes with time-variant covariates and imperfect maintenance effects. *Applied Stochastic Models in Business and Industry*, *37*(3), 592-611. https://doi.org/10.1002/asmb.2600

[30] Zhang, H., Chen, M., Shang, J., Yang, C., & Sun, Y. (2021). Stochastic process-based degradation modeling and RUL prediction: from Brownian motion to fractional Brownian motion. *Science China Information Sciences*, *64*(7), 171201. https://doi.org/10.1007/s11432-020-3134-8

APA Citation

Zhao, X. (2026). Integrated Multi-Sensor Data for Urban Infrastructure Health Monitoring Based on Digital Twin Structural Degradation Modeling. *Journal of Digital Frontier, 1*(2), 58-91. **https://doi.org/10.67541/JDF2610**