

Multi-resolution digital twins under uncertainty with applications to multiscale systems

Stefanos Pyrialakos^{*1}, Anirban Chaudhuri^{†1}, and Karen Willcox^{‡1}

¹*University of Texas at Austin, Austin, TX 78712, United States*

Many physical systems of engineering interest exhibit behavior governed by interactions across various spatial and temporal scales. Accurately capturing these multiscale phenomena in a digital twin requires models of sufficiently high-resolution, yet such models are often too computationally expensive for real-time state estimation, uncertainty quantification, and decision-making. This work addresses this challenge by introducing a multi-resolution digital twin framework that explicitly represents and manages uncertainty while dynamically alternating between low- and high-resolution model representations. The architecture leverages a computationally inexpensive single-scale model that is used for frequent data assimilation and rapid updates, whereas a detailed multiscale finite element model that is invoked selectively to refine predictions when elevated uncertainty or critical events warrant additional resolution. This selective activation enables the digital twin to preserve predictive accuracy without incurring the prohibitive cost of running the high-resolution model continuously. As a proof-of-concept, the proposed framework is demonstrated through a multiscale finite element example for static analysis of an aircraft structural component, where both model resolutions describe the same physical asset but encode different levels of mechanistic detail.

I. Introduction

Digital twins are dynamically updated virtual representations of physical assets that maintain two-way interactions with between the physical assets and their digital counterparts [1–5]. They have emerged as a powerful paradigm for monitoring, prediction, and decision support across aerospace [6–10], civil infrastructure [11–13] and biomedicine [14–16]. Traditional digital twins typically operate at a single model resolution. While computationally tractable, this approach can limit the ability to capture localized phenomena or subtle degradation patterns when only coarse models are available, or conversely, may be prohibitively expensive if only high-fidelity models are used. In general, the physical systems admit multiple meaningful levels of representation, ranging from simplified phenomenological models to high-fidelity physical simulations, each offering different tradeoffs between computational cost and predictive accuracy. Thus, a central challenge is the choice of virtual representation that allows the digital twin to be executed accurately and within realistic timescales. This work develops a probabilistic multi-resolution digital twin capable of coherently integrating and dynamically selecting among, models of different resolutions based on uncertainty, computational cost, and operational needs.

In many engineering contexts, the true modeling need is inherently multi-resolution: the digital twin should be capable of operating on and transitioning among a spectrum of model forms, fidelities, and computational resolutions. These may include reduced-order models, simplified structural approximations, subsystem-based models, high-fidelity finite element (FE) analyses, or full multiscale formulations. A multi-resolution digital twin must therefore reason not only about the latent physical state of the asset but also about which model resolution is most informative at a given time, and how information should move consistently across these resolutions. Developing such capabilities pose several challenges. First, distinct model resolutions do not, in general, share common state spaces or parametrization, which complicates information exchange and consistency. Second, the computational gap across resolutions can be substantial: high-fidelity models (such as detailed three-dimensional FE analyses or nested FE^2 multiscale simulations) may be orders of magnitude more expensive than their coarse or reduced-order counterparts. Third, the “best” resolution is context dependent as fine resolutions may be necessary when the asset is degrading or when measurements detect anomalous behavior, but coarse resolutions often suffice during nominal operation. Existing research provides several

^{*}Postdoctoral Associate, Oden Institute for Computational Engineering and Sciences, stefpyr@oden.utexas.edu, AIAA Member.

[†]Research Scientist, Oden Institute for Computational Engineering and Sciences, anirbanc@oden.utexas.edu, AIAA Senior Member.

[‡]Director, Oden Institute for Computational Engineering and Sciences, Professor of Aerospace Engineering and Engineering Mechanics, kwillcox@oden.utexas.edu, AIAA Fellow

partial ingredients for addressing these challenges. Multifidelity methods offer principled ways to combine models of varying accuracy and cost [17]. Reduced-order and surrogate models are widely used in structural dynamics and SHM to accelerate real-time updating [18–20]. Probabilistic graphical models and Bayesian filters have been used for condition assessment, damage diagnosis, and prediction [8, 21–23]. However, these efforts treat the digital twin as bound to a fixed resolution. Digital twins based on dynamic Bayesian networks (DBNs) [8, 24–26] offer a natural probabilistic structure for the multi-resolution setting.

In this work, we propose a general multi-resolution probabilistic digital twin formulation built on a Probabilistic Graphical Model (PGM) that integrates multiple model resolutions of the same physical asset. The formulation is resolution-agnostic: resolution levels may represent different fidelities, levels of physical abstraction, surrogate versus full-order models, or, when available, multiscale FE formulations. As a case study, we instantiate this general framework with a two-level resolution hierarchy for a composite panel segment under cyclic loading: a fast single-scale finite element model and a high-fidelity FE^2 multiscale model [27, 28] in which mesoscale representative volume elements (RVEs) are solved at macroscale integration points. The choice of resolution is treated as a digital action within a decision-theoretic component of the digital twin, enabling selective escalation to FE^2 analysis only when it materially improves inference. State estimation is performed with an unscented Kalman filter [29], which fuses SHM measurements with predictions issued at whichever resolution is active. This application demonstrates that the proposed multi-resolution digital twin achieves accuracy gains relative to a pure single-scale twin while avoiding the prohibitive cost of continuous multiscale analysis. Sparse multi-resolution updates are sufficient to anchor the digital state, suppress large deviations from the physical state, and capture sudden degradation events, illustrating the promise of multi-resolution frameworks that manage the complexity of the models to balance accuracy and computational cost.

The remainder of the paper is organized as follows. Section II presents the PGM formulation and details how multiple model resolutions are embedded within a unified probabilistic structure. Section III applies the framework to the composite panel SHM case study and evaluates the effect of selective resolution escalation. Section IV summarizes the findings and outlines future research directions.

II. Digital twins for multi-resolution systems

This section describes the multi-resolution problem setup for digital twins followed by a description of the single-resolution PGM. We then describe the proposed multi-resolution digital twins using PGMs and the details of the Bayesian calibration and resolution mapping required for the execution of the digital twins.

A. Multi-resolution digital twin problem formulation

We consider a digital asset represented by a computational model that monitors the physical states of its real-world counterpart for structural systems. Many structural systems exhibit behavior governed by phenomena operating across multiple spatial and temporal scales. To represent these systems within a digital twin, we consider a family of computational models that describe the same asset at different resolutions. Each model emphasizes a distinct balance between physical detail and computational tractability. High-resolution multiscale models can resolve microscale degradation or nonlinear mechanisms, while reduced single-scale models provide simplified dynamics suitable for rapid computation. In this setting, the real-world system is described by *physical state* variables S_t at time t . Let $m = 1, \dots, k$, denote the specific model resolutions used in the digital twins that can represent the physical twin to the desired level. Then the *digital state* is the parameterization of the k computational models denoted by D_t at time t . Each model m predicts the same class of observable quantities but is characterized by different internal state variables and different computational cost; below we will particularize D_t to each model m and address mappings between levels. The digital twin receives *observations* O_t from the physical asset at discrete times t and must assimilate them to infer the latent structural state. These observations may originate from strain sensors, accelerometers, or other structural-health monitoring systems. The digital twin analyzes the performance of the system using *quantities of interest* (QoIs), denoted by Q_t , which are computational estimates of often unobservable physical asset characteristics, such as stress and strain fields.

Since no single-resolution is uniformly optimal for balancing accuracy and cost, the digital twin must determine not only the appropriate *physical action* U_t (e.g., maintenance, inspection) but also a *digital action* $Z_t \in \{1, \dots, k\}$ indicating which model resolution to use at time t . Resolution selection is driven by the digital twin's objective to maintain predictive accuracy while respecting real-time computational constraints. The need for multi-resolution operation introduces an additional requirement: when the chosen resolution changes, the underlying state representation

changes as well. Although all models describe the same physical asset, they employ different digital state representations denoted by $\mathbf{D}_t^{(m)}$ at the specific resolution m . To maintain coherence in the digital twin's current estimate of the structural state, transitions between resolutions are handled through mappings that translate latent digital states from one model's representation to another. These mappings allow the digital twin to switch models without losing accumulated information or introducing inconsistencies. The resolution-agnostic reference digital state obtained after mapping is denoted by \mathbf{D}_t^* , which is shared across all models. There is considerable flexibility in choosing the reference digital state \mathbf{D}_t^* . The optimal form of \mathbf{D}_t^* is problem dependent, and it is ultimately the modeler's decision. For example, if \mathbf{D}_t^* is intended to be used in many-query analyses (e.g., uncertainty propagation), it is generally more practical to define the reference state in the low-resolution space. The augmented digital state representation maintained by a multi-resolution digital twin at time t is given by

$$\mathbf{D}_t = \left[\mathbf{D}_t^*, \mathbf{D}_t^{(1)}, \dots, \mathbf{D}_t^{(k)} \right]^\top. \quad (1)$$

In this paper, we develop the multi-resolution digital twin within a multiscale paradigm consisting of $k = 2$ model resolutions, noting that the ideas extend to more than two levels. Multiscale engineering systems admit multiple computational representations, each resolving the underlying physics at a different level of detail. The coarse resolution ($m = 1$), represented by a single-scale model, provides rapid evaluations based on simplified constitutive behavior. In contrast, the high-resolution ($m = 2$), represented by a multiscale model, resolves additional physics, such as local microstructural phenomena or mesoscale interactions, at a higher computational cost. The single-scale resolution digital state $\mathbf{D}_t^{(1)}$ contains effective macroscale descriptors that parameterize the coarse constitutive behavior of the structure. These descriptors offer a computationally inexpensive prediction model by summarizing the system's response through a small set of effective properties, but they do not encode microscale physics. The multiscale resolution digital state $\mathbf{D}_t^{(2)}$ contains micromechanical descriptors, such as local damage measures, evolving microstructural parameters, and internal variables associated with fine-scale constitutive behavior. The multiscale resolution can represent a significantly richer range of physical phenomena and therefore exhibits stronger predictive capabilities. Since the single-scale resolution is decoupled from microstructural behavior, the reference state \mathbf{D}_t^* must be defined at the macroscopic structural scale. Consequently, after obtaining \mathbf{D}_t^* one can directly run the cheap-to-evaluate macroscale simulation as an approximation for the forward uncertainty propagation. A natural choice in the multiscale structural analysis context is the macroscopic stiffness tensor \mathbf{C}_t as both resolutions can be directly related to it, but other forms of the reference state can be straightforwardly applied.

B. Background: Digital twins using probabilistic graphical models

The multi-resolution digital twin developed in this work builds on the probabilistic graphical model (PGM) formulation established for single-resolution digital twins [8]. In the PGM formulation, the digital twin is represented as a structured probabilistic model in which the digital state, the observed structural response, and the chosen actions are combined into a coherent Bayesian update-and-evaluate cycle. Conditional independencies encoded in the graph enable scalable inference and principled uncertainty quantification, which are essential for real-time structural monitoring. The joint probability distribution that describes the digital twin conditioned on observed data $\mathbf{O}_t = \mathbf{o}_t$ from the physical asset and physical control actions selected by the decision maker $\mathbf{U}_t = \mathbf{u}_t$, for all time till the current time, $t = 0, \dots, t_c$, is formulated as [8]

$$p(\mathbf{D}_0, \dots, \mathbf{D}_{t_c}, \mathbf{Q}_0, \dots, \mathbf{Q}_{t_c} \mid \mathbf{o}_0, \dots, \mathbf{o}_{t_c}, \mathbf{u}_0, \dots, \mathbf{u}_{t_c}) = \prod_{t=0}^{t_c} \left[\phi_t^{\text{update}} \phi_t^{\text{QoI}} \right], \quad (2)$$

where

$$\phi_t^{\text{update}} = p(\mathbf{D}_t \mid \mathbf{D}_{t-1}, \mathbf{U}_{t-1} = \mathbf{u}_{t-1}, \mathbf{O}_t = \mathbf{o}_t), \quad (3)$$

$$\phi_t^{\text{QoI}} = p(\mathbf{Q}_t \mid \mathbf{D}_t), \quad (4)$$

where the two factors, ϕ_t^{update} and ϕ_t^{QoI} , capture the distinct roles of estimating the states and QoIs, respectively. The update factor, ϕ_t^{update} , encodes the assimilation of new data through Bayesian inference. This term couples the process model governing how the digital state evolves with the likelihood model that evaluates the compatibility of observations with predicted responses. The update stage absorbs model-form uncertainty, measurement noise, and nonlinearities present in the state evolution. The factor, ϕ_t^{QoI} , involves forward uncertainty propagation through the computational model to estimate the distribution of the QoIs.

C. Multi-resolution digital twins using probabilistic graphical models

The PGM architecture provides two properties that are essential for extending digital twins to multi-resolution settings. First, the explicit factorization into update and QoI terms clearly separates the data assimilation and QoIs characterizing the physical asset for decision-making. Second, the PGM formulation is modular, such that additional variables, alternative observation models, or different action types can be inserted without restructuring the entire graph. These features make the PGM an ideal foundation for introducing resolution selection as a digital action and for incorporating multiple modeling resolutions within a unified probabilistic framework. To accommodate multi-resolution modeling, we will retain the update–QoI factorization from the graph but condition both factors on the resolution selected at each time step. The resulting hierarchical structure as shown in Figure 1 allows the digital twin to choose between model resolutions as part of its decision process while maintaining probabilistic consistency across model transitions. Figure 1 illustrates the PGM for a specific instance of a multi-resolution digital twin using two resolutions. As shown in Figure 1a, two decisions are made at each timestep, both influencing the digital state D_t . The physical action U_t governs real-world interventions, such as maintenance and repairs, that modify the physical and digital states. The digital action Z_t determines the model resolution used for data assimilation and QoI prediction, and thus directly affects the fidelity of the updated digital state. Figure 1b presents a representative multi-resolution workflow that integrates a single-scale model with a multiscale alternative within a unified digital twin framework to map to the multi-resolution augmented digital state given by Eq. (1).

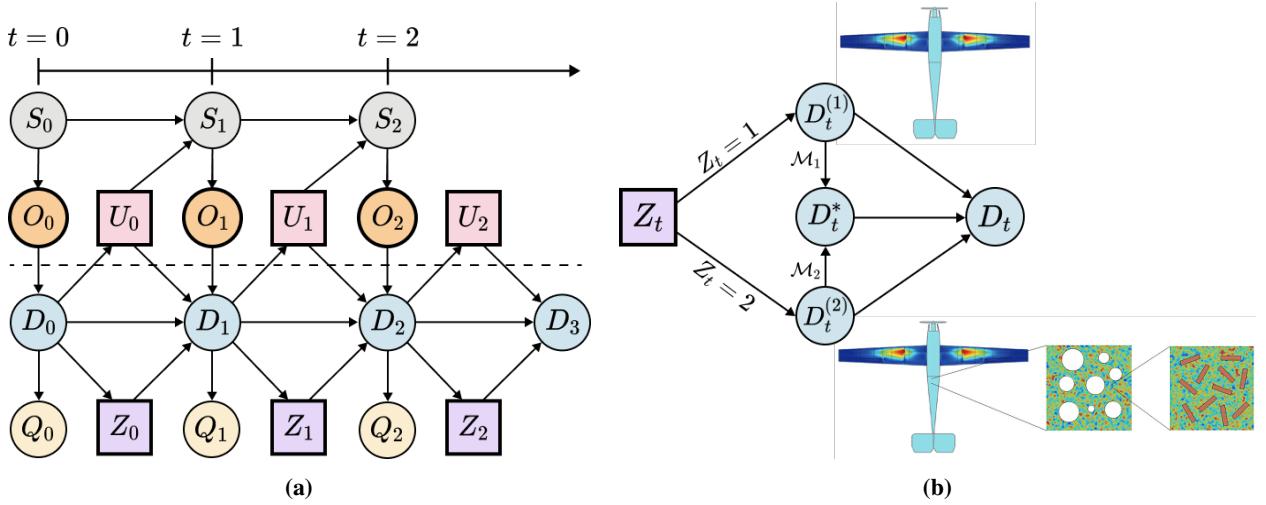


Fig. 1 (a) The probabilistic graphical model (PGM) underlying the multi-resolution digital twin. The digital twin operates across multiple model resolutions, with the digital action Z_t selecting the active resolution and thereby activating its corresponding digital state representation $D_t^{(Z_t)}$. (b) The multi-resolution digital twin explored here spans from a single-scale FE model at its low resolution to a full multiscale FE² model at its high-resolution. When the digital action Z_t picks a certain resolution, the corresponding resolution-specific digital state $D_t^{(Z_t)}$ is activated and then mapped onto a common reference state D_t^* . The augmented digital state representation of the multi-resolution system is denoted by D_t .

The joint probability of the digital states and the QoIs can be represented by the same factorization as in Eq. (2). The update and QoI factors are now conditioned on the digital action Z_t to account for the switching between the various resolution models as given by

$$\phi_t^{\text{update}} = p \left(D_t^*, D_t^{(Z_{t-1})} \mid D_{t-1}, O_t = o_t, U_{t-1} = u_{t-1}, Z_{t-1} = z_{t-1} \right), \quad (5)$$

$$\phi_t^{\text{QoI}} = p \left(Q_t \mid D_t^{(Z_{t-1})}, Z_{t-1} = z_{t-1} \right), \quad (6)$$

with the computational model taken from the resolution indicated by Z_t . We provide details on the different escalation strategies used to define Z_t for the multiscale system in this work in Section II.D. Lower-cost resolutions may use simplified operators, while higher-cost resolutions may employ refined discretizations, richer constitutive laws, or multiscale couplings. Data assimilation (e.g., filtering) proceeds exactly as in the single-resolution case, with Z_t

determining which model is evaluated inside the Bayesian inference step. In this work, we use the unscented Kalman filter (UKF), which performs Bayesian filtering for nonlinear systems using a deterministic set of sigma points that capture the first two moments of the probability distribution [29]. The QoI estimation requires a forward uncertainty propagation through the model selected through Z_t . However, estimating the QoI factor in Eq. (6) can be computationally prohibitive when the active resolution is an expensive-to-evaluate high-resolution model. We define a computationally efficient approximate estimator of the QoI factor, $\hat{\phi}_t^{\text{QoI}}$, by fixing the model resolution used to run the forward uncertainty propagation for QoI estimation to utilize the lower-resolution model capable of using the reference digital state \mathbf{D}_t^* . Note that \mathbf{D}_t^* is selected to be shared across all models as described in Section II.A and maintains a correction mapping from the higher-resolution models as described in Section II.E.

To maintain coherence while switching between various resolution models through the digital action Z_t in the multi-resolution digital twin, we require a mapping function to transform the resolution-specific digital state to the resolution-agnostic reference digital state \mathbf{D}_t^* . We introduce a generic mapping operator

$$\mathcal{M} : \mathbf{D}_t^{(Z_t)}, Z_t \mapsto \mathbf{D}_t^*,$$

which transfers the current digital state $\mathbf{D}_t^{(Z_t)}$ from the active resolution defined by Z_t at time t to the reference digital state \mathbf{D}_t^* . This operator can represent, for example, homogenization, projection, compression to low-order features, or identity if the reference digital state equals a specific resolution state. The specific choice of \mathcal{M} is application-dependent. The PGM structure treats it as a modular component making it easy to define and embed application-specific mapping operators. A detailed definition of the mapping operator for our multiscale problem setup is provided in Section II.E.

D. Multi-resolution escalation strategies for multiscale systems

The multi-resolution character of the digital twin arises from its ability to switch between a low-fidelity single-scale model and a high-fidelity multiscale model. This choice is determined by a digital action Z_t , which governs whether the digital twin operates with only the latent macroscale stiffness parameters or swaps them with multiscale resolution-specific degradation variables by escalating to an FE² computation. In this work, Z_t follows a simple predetermined cadence (e.g., a high-resolution update every five timesteps) chosen to clearly illustrate the effect of switching between resolutions. Extensions toward uncertainty-driven or value-of-information policies are natural within this framework and are currently under investigation. We compare the following three escalation strategies:

- 1) a purely single-scale baseline, in which the multiscale model is never used and the filtering algorithm relies exclusively on the macroscale stiffness factors to explain the observations;
- 2) an always-multiscale configuration, in which the multiscale model is used at every timestep and the macroscale response is purely determined by the evolving micromechanical material parameters;
- 3) a hybrid configuration, in which the multiscale model is activated every five timesteps and the single-scale model is used at all other times.

The first case corresponds to a single-resolution digital twin with no micro-structural resolution. The second case provides a reference solution in which the microscopic damage state is tightly coupled to the observed response but incurs the highest computational cost. The third case aims to approximate the accuracy of the full multiscale solution at significantly reduced cost by using the multiscale corrections sparsely in time.

E. Multi-resolution mapping between digital states for multiscale systems

In the multi-resolution setup for multiscale systems with two model resolutions, we choose the macroscopic stiffness tensor \mathbf{C}_t as the resolution-agnostic reference digital state \mathbf{D}_t^* . The constitutive stiffness tensor \mathbf{C}_t has the following general form:

$$\mathbf{C}_t = \begin{bmatrix} C_{11,t} & C_{12,t} & C_{13,t} & 0 & 0 & 0 \\ & C_{22,t} & C_{23,t} & 0 & 0 & 0 \\ & & C_{33,t} & 0 & 0 & 0 \\ & & & C_{44,t} & 0 & 0 \\ & & & & C_{55,t} & 0 \\ & & & & & C_{66,t} \end{bmatrix}, \quad (7)$$

where each non-zero component $C_{ij,t}$, for $i, j = 1, \dots, 6$, is assumed to evolve independently of the others. The constitutive tensor \mathbf{C}_t is the linear operator that maps strain to stress in a material. At a certain timestep t , each

component C_{ij} , for $i, j = 1, \dots, 6$, represents a specific stiffness contribution, capturing normal stiffness (C_{11}, C_{22}, C_{33}), shear stiffness (C_{44}, C_{55}, C_{66}) or the coupling between different deformation directions (C_{12}, C_{13}, C_{23}).

Each resolution provides a predicted macroscopic stiffness through a resolution-specific mapping operator as

$$\mathbf{D}_t^* = \mathcal{M}(\mathbf{D}_t^{(Z_t)}, Z_t) = \begin{cases} \mathcal{M}_1(\mathbf{D}_t^{(1)}), & Z_t = 1 \\ \mathcal{M}_2(\mathbf{D}_t^{(2)}), & Z_t = 2. \end{cases} \quad (8)$$

For the single-scale resolution, the operator \mathcal{M}_1 maps a set of effective macroscale descriptors to the macroscopic stiffness tensor. This operator typically represents reduced-order stiffness parametrizations, parametric constitutive relations, or closed-form effective-medium approximations. It defines a low-cost operation and does not resolve microscale physics. The operator \mathcal{M}_1 has the following form:

$$\mathcal{M}_1(\mathbf{D}_t^{(1)}) = \mathbf{A}_t^\top \mathbf{C}_{\text{healthy}} \mathbf{A}_t + \Delta \mathbf{C}_t, \quad (9)$$

where the matrix $\mathbf{C}_{\text{healthy}}$ is the fully healthy macroscopic stiffness tensor. In this work, the matrix \mathbf{A} applies a transversely isotropic stiffness reduction on $\mathbf{C}_{\text{healthy}}$ and is defined, at timestep t , as:

$$\mathbf{A}_t = \text{diag} \left(\sqrt{a_{L,t}}, \sqrt{a_{L,t}}, \sqrt{a_{T,t}}, \sqrt{a_{T,t}}, \sqrt{a_{S,t}}, \sqrt{a_{S,t}} \right). \quad (10)$$

where $a_{L,t}$, $a_{T,t}$ and $a_{S,t}$, at timestep t , denote the longitudinal, transverse and shear stiffness reduction factors that scale the fully healthy stiffness tensor $\mathbf{C}_{\text{healthy}}$. The term $\Delta \mathbf{C}_t$ is a stiffness correction matrix, whose precise form is defined below in Eq. (12).

For the multiscale resolution, the operator \mathcal{M}_2 maps micromechanical descriptors to the macroscopic stiffness tensor. This mapping is generally referred to as homogenization. Depending on the chosen homogenization strategy, \mathcal{M}_2 may range from analytical formulations to the numerical solution of a discretized boundary-value problem. In this work, we focus on a computational homogenization operator [30], which takes the generic form,

$$\mathcal{M}_2(\mathbf{D}_t^{(2)}) = \frac{1}{|\Omega_{\text{fine}}|} \int_{\Omega_{\text{fine}}} \frac{\partial \boldsymbol{\sigma}_{\text{fine},t}(x; \mathbf{D}_t^{(2)})}{\partial \boldsymbol{\varepsilon}_t} dx, \quad (11)$$

where Ω_{fine} denotes the representative volume element (RVE) associated with the fine length scale (e.g. microscale, mesoscale), x is the spatial coordinate within this domain, $\boldsymbol{\sigma}_{\text{fine},t}(x; \mathbf{D}_t^{(2)})$ is the fine scale stress field response, at timestep t , driven by the fine-scale digital states, and $\boldsymbol{\varepsilon}_t$ is the corresponding macroscale deformation measure.

In a multiscale setting, the equilibrium problem is solved using the FE method, with constitutive data supplied by the solution of a FE-based RVE at each integration point. This procedure defines the well-established FE² algorithm [31], which operates as follows: the current macroscale deformation measure $\boldsymbol{\varepsilon}_t$ is passed to the RVE, the corresponding microscale boundary value problem is solved with $\boldsymbol{\varepsilon}_t$ prescribed as the boundary condition, and the resulting homogenized macroscale Cauchy stress $\boldsymbol{\sigma}_t$ and stiffness tensor \mathbf{C}_t are returned to the macroscale solver. Whenever the digital twin escalates to the multiscale resolution, this two-way coupled problem is solved directly. Although the FE² algorithm provides high-fidelity predictions, its computational cost is prohibitive for use in a digital twin setting.

The mapping operators at both resolutions yield the macroscopic stiffness tensor \mathbf{C}_t as shown in Eq. (8). However, discrepancies may arise between these predictions for several reasons. For example, the coarse-resolution parametrization may impose simplifying assumptions, such as transverse isotropy, while the micromechanical description at the multiscale level may allow the material to evolve into a non-isotropic way, such as orthotropic, during degradation. To preserve coherence between resolutions, whenever there is escalation to the multiscale model, we introduce an additive correction factor:

$$\Delta \mathbf{C}_t = \mathbf{C}_t^{(2)} - \mathbf{C}_t^{(1)}. \quad (12)$$

where $\mathbf{C}_t^{(1)}$ and $\mathbf{C}_t^{(2)}$ represent the macroscopic constitutive stiffness tensor estimated by the single-scale and multiscale resolutions, respectively. The factor in Eq. 12 provides a partition-wise stiffness correction that quantifies the discrepancy between the limited single-scale prediction and the more accurate multiscale prediction. This correction is incorporated into the shared reference stiffness state \mathbf{D}_t^* and carried forward through the following low-resolution steps as:

$$\mathbf{D}_t^* = \mathbf{C}_t = \mathbf{C}_t^{(Z_t)} + \Delta \mathbf{C}_t. \quad (13)$$

The mechanism ensures that information gained from expensive high-fidelity evaluations is retained, stabilizing the macroscopic parameters and suppressing artificial discontinuities when switching between models. When all timesteps use the multiscale model, the corrected latent stiffness coincides with the FE^2 -based stiffness. When only the single-scale model is used, the correction remains inactive and the stiffness evolution is governed solely by the macroscale parameters. For the hybrid case, if the active resolution is single-scale ($Z_t = 1$), the correction injects fine-scale information into the reference state. If the active resolution is multiscale ($Z_t = 2$), then $\Delta\mathbf{C}_t = 0$. Hence, multiscale evaluations provide persistent refinements to the coarse representation.

III. Application: Multi-resolution digital twin for a composite panel

The proposed multi-resolution digital-twin framework is demonstrated on a cantilevered composite-panel segment. The digital twin integrates two modeling resolutions of the same physical asset: a single-scale finite element model serving as the low-resolution representation, and a multiscale FE^2 model providing the high-resolution counterpart.

A. Resolution-dependent model configuration

The macroscale structure represents a typical wing panel and is modeled as a straight cantilever with span of 6 m, chord length of 1 m, and thickness 0.15 m. The domain is discretized by a structured hexahedral mesh with $18 \times 3 \times 2$ elements along the span, chord and thickness directions, respectively. The mesh is partitioned into eight blocks along the span, each block representing a monitoring partition for which a separate degradation state is inferred. Only the partitions 2, 4, and 6 are assumed to have an evolving structural health state, while the rest are taken as fully healthy throughout the analysis. The boundary conditions correspond to a clamped root plane at the inboard end, where all displacements are suppressed on the plane normal to the span direction. A uniform pressure load of magnitude 5×10^4 Pa is applied to the upper surface, representing an aerodynamic lift load. Figure 2 illustrates the macroscale geometry, mesh and the partitioning of the panel. All structural responses in this application are computed under the assumption of quasi-static equilibrium, as inertial effects are negligible for the loads and degradation rates considered. At each update step t , the governing equilibrium equation is:

$$\nabla \cdot \boldsymbol{\sigma}_t + \mathbf{b}_t = \mathbf{0}, \quad (14)$$

where $\boldsymbol{\sigma}_t = \mathbf{C}_t : \boldsymbol{\varepsilon}_t$ is the Cauchy stress tensor, at timestep t , obtained from the constitutive stiffness \mathbf{C}_t , and \mathbf{b}_t denotes the applied body forces at t . This quasi-static assumption applies uniformly to both the single-scale and the FE^2 multiscale evaluations used within the digital-twin framework.

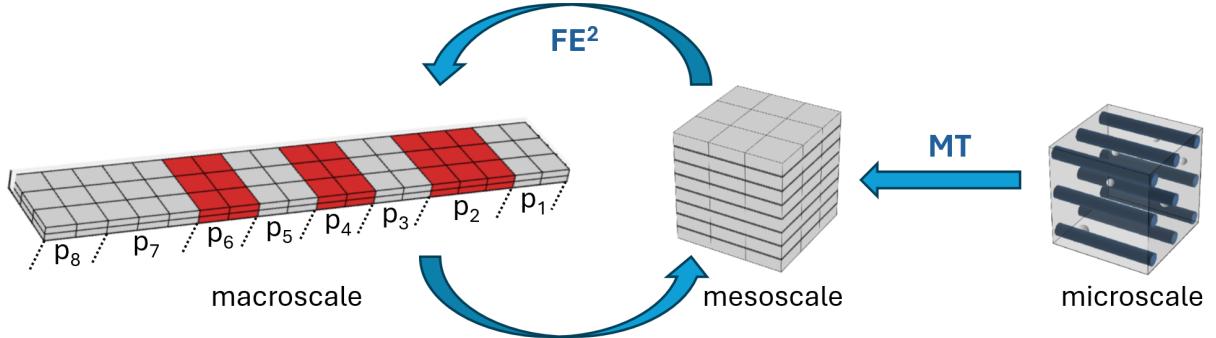


Fig. 2 Multi-resolution representation of the structural component. The macroscale finite element model includes three degrading partitions (partitions 2, 4, and 6 shown in red) whose effective stiffness deteriorates during operation. The single-scale model corresponds to this macroscale level and remains agnostic to meso- and micro-scale behavior. The multiscale model extends the hierarchy with two additional resolutions: a mesoscale FE-RVE that resolves the laminate architecture, including stacked plies and imperfect interply interfaces, and interacts with the macroscale through an FE^2 scheme; and a microscale Mori–Tanaka model that captures constituent behavior, fiber distributions, and the evolution of matrix porosity.

The low-resolution model is a single-scale formulation defined solely at the macroscale, with no explicit coupling to finer length scales. The macroscopic response is characterized through a digital state comprising three multiplicative

stiffness-scaling factors, $\mathbf{D}_t^{(1)} = [\{a_{L,t}^p, a_{T,t}^p, a_{S,t}^p\}_{p=1}^3]$, where $p = 1, 2, 3$ correspond to the three degrading partitions, namely partitions 2, 4, and 6. For each degrading partition, $a_{L,t}^p$, $a_{T,t}^p$ and $a_{S,t}^p$, are the stiffness reduction factors defined in Eq. 10. These parameters evolve continuously and constitute the low-fidelity resolution-specific state. This model preserves computational efficiency by enforcing transverse isotropy while still allowing directional softening as defined in Eq. (9). This model can be used for rapid data assimilation during the digital twin operation; however, its limited predictive capabilities naturally impact the quality of the digital twin.

The high-resolution representation is defined by a multiscale model comprised of three distinct length scales: the microscale, the mesoscale, and the macroscale. At the microscale, the material is modeled as a unidirectional carbon-fiber-reinforced polymer with epoxy matrix, carbon fibers, and voids as distinct phases. Baseline elastic moduli are such that the matrix stiffness is 3.5 GPa, the fiber stiffness 240 GPa, and the fiber volume fraction 60%. Voids are treated as compliant inclusions to preserve positive definiteness of the effective stiffness. A Mori–Tanaka homogenization law [32] provides a closed-form estimate of the composite stiffness (see Appendix A) with a fixed volume fraction of fiber inclusions and an evolving volume fraction of voids, whose evolution is defined by the digital state ψ_t . This estimate ultimately defines the porosity dependent effective stiffness of the individual plies on the mesoscale level. At the mesoscale, individual plies and interply layers are resolved through a finite element RVE that captures interlaminar deterioration, which is described by the digital state $d_{int,t}$, and, inherited from the microscale, the influence of porosity on ply stiffness. The degradation variables form the multiscale digital state as $\mathbf{D}_t^{(2)} = [\{\psi_t^p, d_{int,t}^p\}_{p=1}^3]$, which is active only when the high-resolution model is invoked. The FE² formulation computes homogenized stresses and tangent moduli from these variables at every macroscale integration point of degrading partitions, providing partition-wise effective stiffnesses that directly link these micromechanical parameters to macroscopic behavior. Although this multiscale model has the ability to provide very accurate and more interpretable predictions, the need to perform a computationally intensive FE² simulation within the downstream tasks of data assimilation and forward uncertainty propagation makes it prohibitively expensive.

Across resolutions, the digital twin maintains a persistent latent representation of the nine independent components of the macroscopic stiffness tensor defined in Eq. (7) for each of the three monitored partitions. Resolution-specific digital states are activated only when their corresponding fidelity is in use, and each employs a distinct mapping to and from the shared latent space as described in Section II.E through Eqs. (9) and (11) for the low- and high-resolution models, respectively. The augmented digital state vector for the multi-resolution twin, \mathbf{D}_t , gathers all resolution-specific and resolution-agnostic state variables and is defined as

$$\mathbf{D}_t = \underbrace{\left[\{a_{L,t}^p, a_{T,t}^p, a_{S,t}^p\}_{p=1}^3 \right]}_{\mathbf{D}_t^{(1)}} \cup \underbrace{\left[\{\psi_t^p, d_{int,t}^p\}_{p=1}^3 \right]}_{\mathbf{D}_t^{(2)}} \cup \underbrace{\left[\{C_{11,t}^p, C_{22,t}^p, C_{33,t}^p, C_{12,t}^p, C_{13,t}^p, C_{23,t}^p, C_{44,t}^p, C_{55,t}^p, C_{66,t}^p\}_{p=1}^3 \right]}_{\mathbf{D}_t^*}.$$

This vector persists throughout the simulation; however, at each timestep only the subset corresponding to the currently active resolution undergoes updating, while the remaining components are carried forward unchanged.

B. Synthetic data generation with degradation dynamics

Synthetic “ground truth” health-state histories are generated by evolving the micromechanical degradation variables contained in the fine-resolution digital state vector $\mathbf{D}_t^{(2)}$. These histories are produced under a prescribed sequence of equivalent load cycles applied over a finite time horizon consisting of 50 discrete update steps. Each update step corresponds to an increment in the accumulated number of equivalent load cycles, denoted ΔN_t . This increment is obtained from a usage-mapping that converts operational flight data into cycle accumulation.

The total number of equivalent cycles accrued over h hours of operation is defined in continuous time by

$$N(h) = \int_0^h r(\tau) s(\tau)^c d\tau \quad (15)$$

where h is the elapsed flight time in hours, and τ is the continuous time variable taking values in the interval $[0, h]$. The function $r(\tau)$ denotes the instantaneous cycles-per-hour usage rate, and $s(\tau)$ is the corresponding normalized load-severity measure taking values in $[0, 1]$. The parameter c is a prescribed load-sensitivity exponent that controls the influence of the severity level on cycle accumulation.

In the discrete form of Eq. 15 used for generating synthetic health histories, the time horizon is partitioned into segments indexed by $j = 1, \dots, t$, where t denotes the current discrete timestep. Each segment has duration

Δh_j , satisfying $\sum_{j=1}^t \Delta h_j = h$. Evaluating the usage rate and severity at each segment midpoint yields the discrete accumulation rule:

$$N_t \approx \sum_{j=1}^t \Delta N_j = \sum_{j=1}^t r_j s_j^c \Delta h_j, \quad (16)$$

where r_j and s_j represent the cycles-per-hour rate and normalized severity associated with the j -th flight segment.

The porosity variable ψ_t in each degrading partition is modeled as a bounded, saturating quantity. Following standard microdamage phenomenological laws [33], its evolution as a function of accumulated cycles N_t is taken as

$$\psi_t = \psi_0 + (\psi_{\max} - \psi_0) \left[1 - \exp \left(- \left(\frac{N_t}{N_{\text{sat}}} \right)^{\beta} \right) \right], \quad \beta > 0, \quad (17)$$

where ψ_0 is the initial porosity, ψ_{\max} is its asymptotic upper bound, while N_{sat} and β are empirical parameters governing the curvature and rate of the saturation response.

The interply degradation variable $d_{\text{int},t}$ is defined through an interfacial effectiveness parameter $\eta_{f,t}$, which represents the fraction of the original interfacial stiffness retained at timestep t . Under cyclic loading, $\eta_{f,t}$ decays according to a cohesive-fatigue (Paris-like) law [34] in the normalized energy-release range $\Delta G_t/G_c$, where ΔG_t denotes the cycle energy-release range, at timestep t , in the current block and G_c is the interfacial fracture-toughness parameter. The evolution law is given by

$$\eta_{f,t} = \left[1 + q_f A_f \left(\frac{\Delta G_t}{G_c} \right)^{m_f} N_t \right]^{-1/q_f}. \quad (18)$$

where A_f , m_f and q_f are material parameters that define how fast and how nonlinearly $\eta_{f,t}$ decays. Since $\eta_{f,t} \in [0, 1]$ acts as an interfacial stiffness-retention factor, the interply damage variable is defined as

$$d_{\text{int},t} = \max \{ (1 - \eta_{f,t}), d_{\max} \}, \quad (19)$$

with d_{\max} describing the limiting interfacial degradation as N_t increases.

The synthetic health-state trajectories obtained using Eqs. (17) and (18) are used to generate strain histories by solving the forward multiscale model. We assume a set of 24 virtual strain gauges distributed along the physical asset, capable of measuring normal strains in the span, chord, and thickness directions within each partition. Specifically, each partition has a set of 3 strain gauges, each monitoring the strain magnitude in a different direction. At each timestep, the multiscale model supplies the partition-wise stiffness tensors, which are passed to the macroscale finite element solver to compute the strain response at the gauge locations under the prescribed pressure loading. The resulting strain signals are then contaminated with independent, zero-mean additive Gaussian noise with standard deviation set to 2% of the local strain magnitude.

We consider three representative degradation scenarios by selectively modifying the material parameters in Eqs. (17) and (18). First, a slow, nearly linear evolution of both porosity and interply degradation is produced by a low-severity, smooth configuration for the interfacial law and for the porosity kinetics. Second, a highly nonlinear but still smooth degradation is generated using a more aggressive cohesive-fatigue response and fast porosity kinetics; this yields rapid early-stage deterioration followed by a progressively slower approach to the asymptotes. Finally, a regime with sharp jumps superimposed on a moderately nonlinear smooth evolution, where two interior timesteps are selected for each variable ψ_t and $d_{\text{int},t}$ and, at each selected time, the corresponding variable is instantaneously increased by 10.0% in relation to its asymptotic limit. Partitions designated as healthy remain fixed at their initial undamaged state throughout. We provide details for all three scenarios in Table 1.

Table 1 Parameter sets for the three representative degradation scenarios.

Scenario	$\Delta G/G_c$	A_f	m_f	q_f	N_{sat}	β	d_{\max}	ψ_{\max}	Jumps
Slow, nearly linear, smooth	0.25	5×10^{-3}	1.0	3.0	10.0	1.0	0.60	0.20	Disabled
Sharp, highly nonlinear, smooth	0.30	3×10^{-1}	2.5	0.4	0.5	3.0	0.60	0.20	Disabled
Sudden-jumps, moderately nonlinear	0.30	1.5×10^{-1}	1.5	0.6	1.5	1.3	0.60	0.20	Enabled

C. Bayesian state estimation implementation details

The Bayesian state estimation employs an Unscented Kalman Filter (UKF) on the augmented multi-resolution digital state $\mathbf{D}_t = (\mathbf{D}_t^*, \mathbf{D}_t^{(1)}, \mathbf{D}_t^{(2)})$, with the active resolution selected by the digital action $Z_t \in \{1, 2\}$. The resolution-specific digital state evolves as

$$\mathbf{D}_t = f^{(Z_t)}(\mathbf{D}_{t-1}) + w_t^{(Z_t)}, \quad (20)$$

$$\mathbf{O}_t = h^{(Z_t)}(\mathbf{D}_t) + v_t^{(Z_t)}, \quad (21)$$

where Eqs. (20) and (21) represent the prediction and data assimilation contributions of the stochastic filtering problem, respectively. The functions $f^{(Z_t)}$ and $h^{(Z_t)}$ are the resolution-specific process and observation models. The observation operator incorporates the stiffness prediction $\mathbf{C}_t^{(Z_t)}$ into the appropriate solver. At the coarse resolution, $f^{(1)}$ governs effective macroscale properties and $h^{(1)}$ evaluates a low-cost model using $\mathbf{C}_t^{(1)}$. At the fine resolution, $f^{(2)}$ updates micromechanical descriptors and $h^{(2)}$ uses the homogenized stiffness $\mathbf{C}_t^{(2)}$ within a higher-fidelity solver. Although the degradation laws are known from the synthetic data generation, we deliberately avoid prescribing them and instead implement both $f^{(1)}$ and $f^{(2)}$ as identity mappings with additive Gaussian noise, placing the burden of learning temporal evolution on the observations while maintaining smooth trajectories. Noise levels match those used in the synthetic experiments: 2×10^{-3} for porosity, 4×10^{-3} for inter-ply damage, and 2×10^{-3} for macroscopic stiffness parameters. State components of inactive resolutions evolve only under process noise. During prediction, sigma points from the current posterior are propagated through $f^{(Z_t)}$.

The observation step converts each propagated sigma point into a strain response via the resolution-dependent operator $h^{(Z_t)}$. Under the single-scale configuration ($Z_t = 1$), the predicted macroscopic stiffness $\mathbf{C}_t^{(1)}$ is obtained from the coarse mapping operator $\mathcal{M}_1(\mathbf{D}_t^{(1)})$ and passed to the macroscale finite element solver, which computes the corresponding strain predictions. When the multiscale resolution is active ($Z_t = 2$), the micromechanical descriptors determine the homogenized stiffness through the FE²-based mapping $\mathcal{M}_2(\mathbf{D}_t^{(2)})$, requiring an RVE solve for each sigma point before the macroscale response is evaluated. In both cases, the observation noise used in the UKF update matches the heteroscedastic noise model of the synthetic measurement process, ensuring a consistent assimilation framework.

Whenever the system escalates to the multiscale resolution, the stiffness mismatch in Eq. (12) is injected into the shared reference stiffness state \mathbf{D}_t^* , implementing the correction mechanism described in Eq. (13). This accumulation ensures that information learned during multiscale steps persists during subsequent single-scale updates, thereby preventing discontinuities in the stiffness trajectory when switching between resolutions. When $Z_t = 2$, the correction vanishes, so repeated multiscale steps recover the FE²-consistent stiffness exactly.

The UKF is controlled by three parameters, namely $\alpha = 0.7$ which controls the spread of the sigma points, β which incorporates any prior knowledge of the distribution, and κ is a secondary scaling parameter. In this study, the parameters were configured as $\alpha = 0.7$, $\beta = 2.0$, and $\kappa = -5$ to regulate the spread of sigma points. A small covariance inflation of $10^{-4}I$ is applied at each iteration to counteract artificial covariance collapse associated with repeated nonlinear transformations and to maintain numerical robustness. A post-update projection is employed when the high-resolution model is active to map posterior uncertainty in the micromechanical digital states to an equivalent uncertainty in the macroscale stiffness states using an unscented transform.

D. Results

We evaluate the three escalation strategies defined in Section II.D under three representative degradation scenarios described in Section III.B. In all cases, stiffness correction is enabled to stabilize the macroscale parameters between high-resolution steps. We show the degradation with respect to the flight hours for the degrading partitions (here, partitions 2, 4, and 6). The degradation is defined as the stiffness ratio $\frac{\|\mathbf{C}_t\|}{\|\mathbf{C}_{\text{healthy}}\|}$, where $\|\mathbf{C}_t\|$ and $\|\mathbf{C}_{\text{healthy}}\|$ are the Frobenius norms of the current state of the stiffness matrix and the fully healthy state of the stiffness matrix, respectively. We present the mean predictions and the uncertainty in the predictions through a $\pm 2SD$ interval, where SD is the standard deviation, to approximately represent the 95% confidence interval.

Figure 3 shows the trajectories of the stiffness ratio for the slow degradation scenario. The hybrid strategy closely follows the always-high-resolution reference throughout the time horizon, with only minor deviations between high-resolution updates. The low-resolution strategy, in contrast, accumulates noticeable bias due to its limited sensitivity to the evolving degradation variables. Figure 4 shows the corresponding stiffness ratio relative RMSE from the prediction mean. Both the hybrid and the always-high-resolution strategies maintain a low relative RMSE, demonstrating that

sparse high-resolution updates are sufficient when degradation evolves smoothly, whereas the low-resolution strategy exhibits consistently higher error.

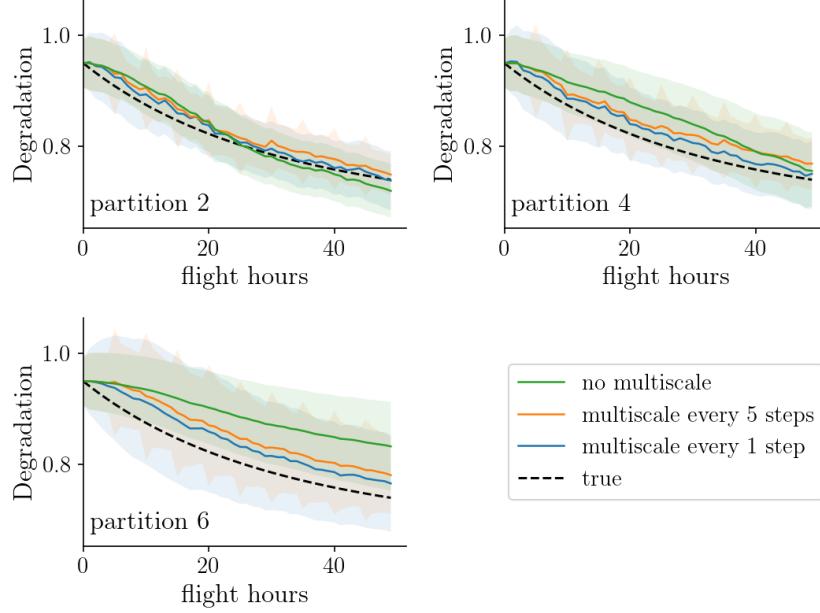


Fig. 3 Estimated stiffness ratios in the slow degradation scenario. The hybrid strategy maintains trajectories close to the always-high-resolution reference while requiring fewer high-resolution updates. Solid lines represent the mean and shaded bands denote $\pm 2SD$ uncertainty.

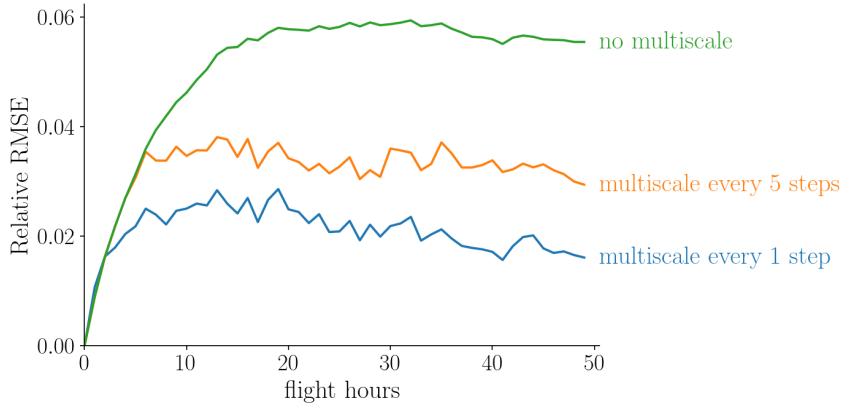


Fig. 4 Macroscopic stiffness ratio error in the slow degradation scenario. The root-mean-square error remains low for both the always-high-resolution and hybrid strategies, demonstrating the effectiveness of sparse high-resolution updates when degradation evolves smoothly.

The sharp degradation scenario highlights behavior during rapid stiffness loss. Figure 5 shows that the always-high-resolution strategy most accurately tracks the steep decline, while the hybrid strategy exhibits modest deviations between escalation points for partitions 4 and 6. Nevertheless, the hybrid strategy preserves the major features of the stiffness trajectory, whereas the low-resolution strategy underestimates the severity of degradation. The RMSE behavior in Figure 6 shows the same trend. The always-high-resolution strategy yields the lowest errors, while the hybrid strategy

achieves an improvement over the low-resolution baseline, with temporary increases occurring only between escalation steps.

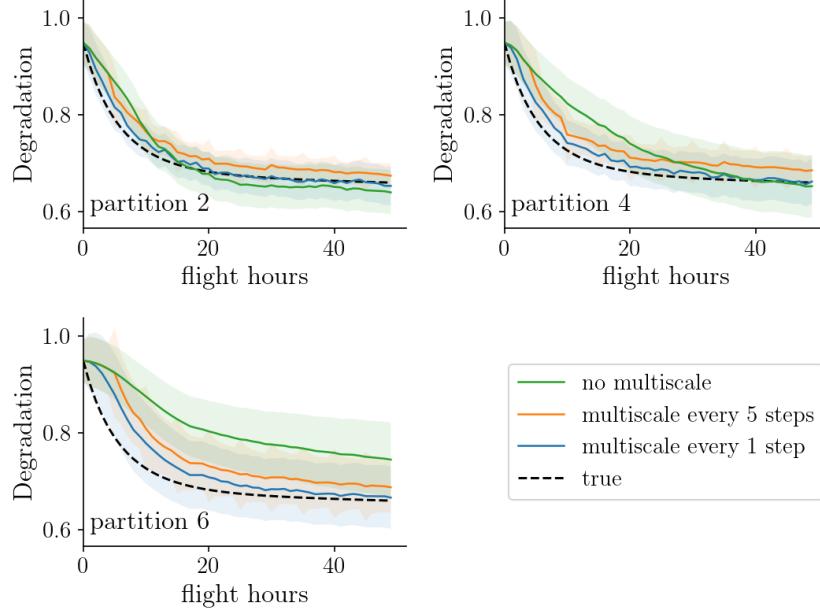


Fig. 5 Estimated stiffness ratios in the sharp degradation scenario. Rapid stiffness loss is captured most accurately by the always-high-resolution strategy; the hybrid strategy tracks this behavior with modest deviations between high-resolution updates. Solid lines represent the mean and shaded bands denote $\pm 2SD$ uncertainty.

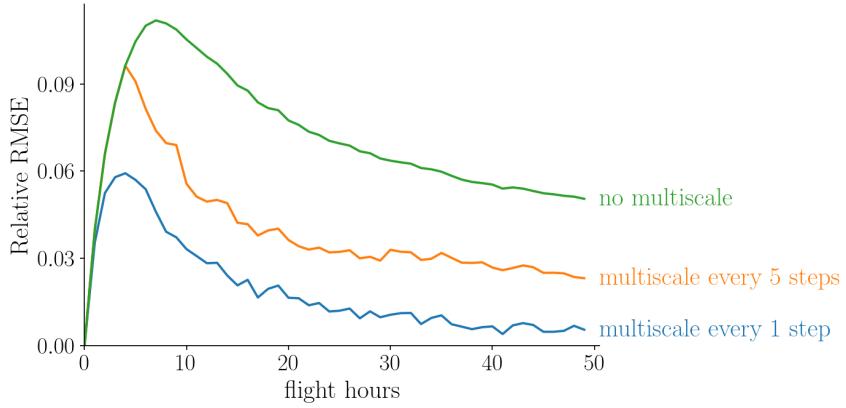


Fig. 6 Macroscopic stiffness ratio error in the sharp degradation scenario. Larger excursions occur during intervals without high-resolution updates, but the hybrid strategy remains significantly more accurate than the purely low-resolution approach.

The sudden-jump degradation scenario tests the ability of each strategy to handle abrupt changes in stiffness. As shown in Figure 7, the hybrid strategy effectively locks in each abrupt stiffness drop identified by the high-resolution model. In contrast, the low-resolution strategy smooths out these jumps and fails to recover their magnitude. The relative RMSE plot in Figure 8 highlights this effect: high-resolution steps sharply reduce prediction error immediately following each jump, and the hybrid strategy retains lower relative RMSE than the low-resolution baseline throughout.

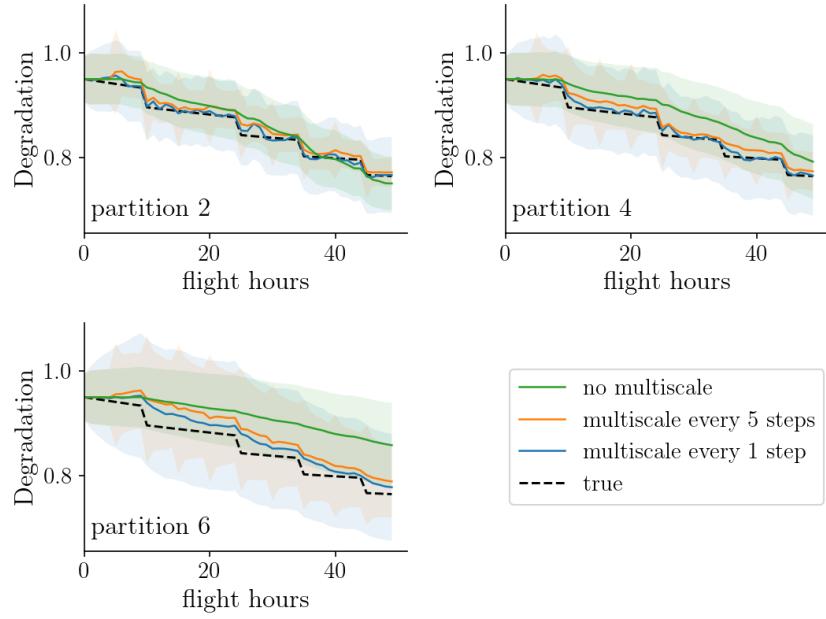


Fig. 7 Estimated stiffness ratios in the sudden-jump degradation scenario. Each abrupt degradation event is locked in by the high-resolution model, and the correction mechanism ensures that the hybrid strategy preserves these stiffness drops during low-resolution intervals. Solid lines represent the mean and shaded bands denote $\pm 2SD$ uncertainty.

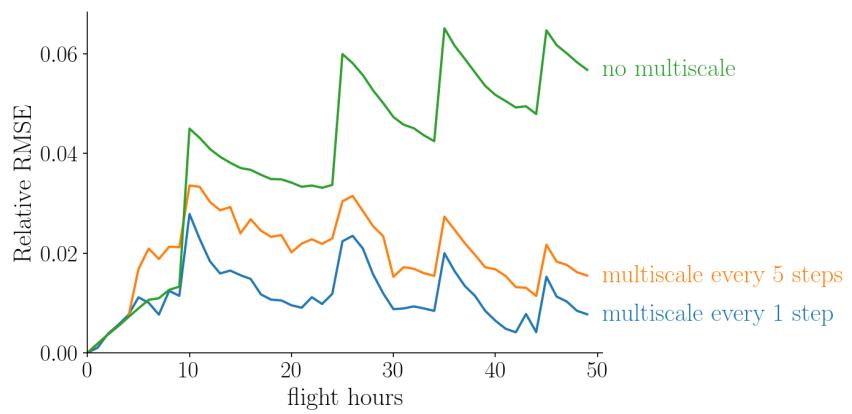


Fig. 8 Macroscopic stiffness ratio error in the sudden-jump degradation scenario. High-resolution steps sharply reduce the relative RMSE immediately after each jump, while the hybrid strategy maintains lower error than the low-resolution baseline.

Next, we examine how the different escalation strategies influence the predicted distribution of a structural QoI at selected monitoring times. For the sudden-jumps degradation scenario, and for each escalation strategy, we perform forward uncertainty–propagation by sampling from the reference digital states D_t^* at three representative flight hours (early: $t = 10$, mid-life: $t = 25$, and late: $t = 50$). For each sampled reference state, we compute the macroscale response using the corresponding stiffness matrix and record the vertical displacement at several locations along the longitudinal direction of the panel. Using 500 Monte Carlo samples, we then construct histograms of the tip displacement. “True” vertical displacements are obtained by solving the forward multiscale problem with the true digital states produced during the synthetic-data generation. Figure 9 presents the resulting displacement distributions: rows correspond to flight hours, and columns to positions along the span. Each subplot overlays the results from the low-resolution, hybrid, and always-high-resolution strategies. The figures show that the modes of the always-high-resolution and hybrid histograms generally fall closest to the true displacement, with the former distribution being naturally the most accurate. Conversely, the always-low-resolution histograms indicate a consistent underestimation of the displacement. The relative RMSE between the mean values of the true and estimated displacement, averaged over spatial positions, for the high-resolution, hybrid, and low-resolution strategies is: for flight hour = 10, 0.28×10^{-3} , 1.85×10^{-3} , and 9.43×10^{-3} ; for flight hour = 25, 0.19×10^{-3} , 1.75×10^{-3} , and 9.92×10^{-3} ; and for flight hour = 50, 0.56×10^{-3} , 1.01×10^{-3} , and 11.30×10^{-3} , respectively.

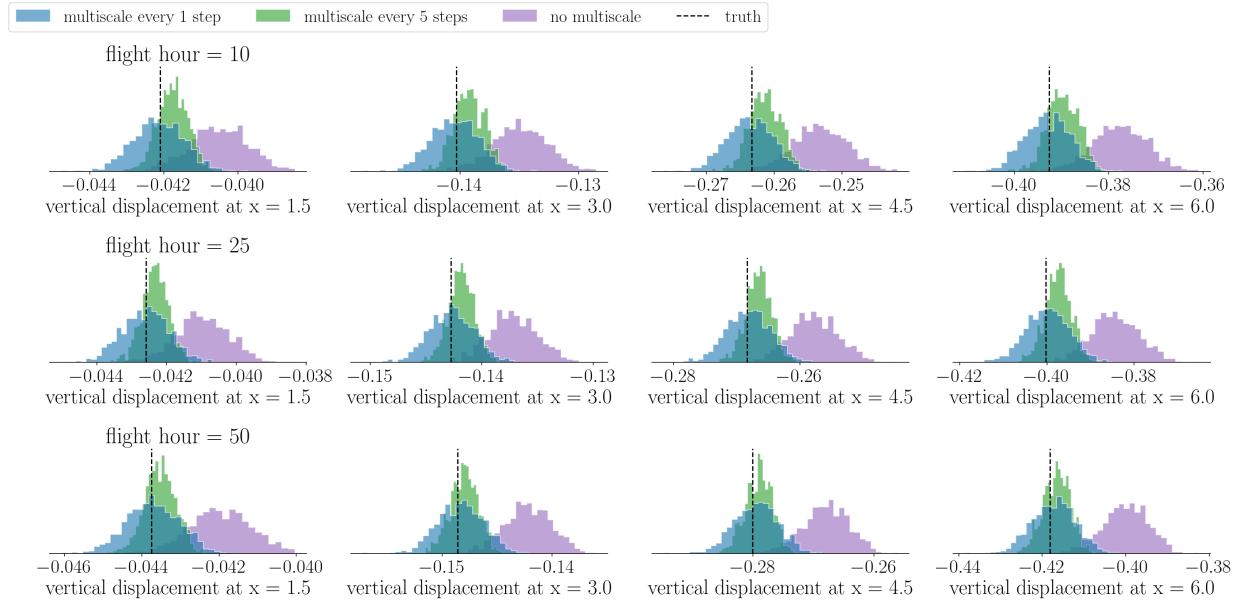


Fig. 9 Displacement histograms in the sudden-jump degradation scenario. The predicted vertical displacement is compared for the three different escalation strategies across several longitudinal positions and flight hours.

Table 2 summarizes the computational cost over a time horizon of 50 steps. The cost ratio is expressed relative to the purely low-resolution baseline. The always-high-resolution strategy is approximately three orders of magnitude more expensive than the low-resolution baseline, reflecting the cost of FE^2 -based homogenization. The hybrid strategy reduces this cost by approximately a factor of five while preserving most of the predictive benefit of continuous high-resolution operation. Across all degradation regimes, the hybrid strategy achieves a near-optimal balance between accuracy and cost. Sparse high-resolution updates, combined with stiffness correction, preserve fidelity during periods of rapid degradation while keeping the overall computational burden compatible with online digital twin operation.

IV. Conclusions

This paper presents a probabilistic framework for multi-resolution digital twins capable of operating across multiple model resolutions within a unified PGM. The approach extends the single-resolution PGM formulation used in conventional digital twins by introducing a digital action and a resolution-agnostic reference digital state to the PGM. The digital action allows the multi-resolution digital twin to dynamically switch between coarse- and high-resolution

Table 2 Computational cost comparison for the three escalation strategies

Strategy	High-resolution steps	Low-resolution steps	Cost ratio
Low-resolution only	0	50	1
Always high-resolution	50	0	~ 1125
Hybrid (every 5 steps)	10	40	~ 225

models during the execution of the digital twin to balance accuracy and computational cost. The reference digital state maintains a coherent reference representation across all resolutions and is obtained through a mapping operator that ensures information gained at one resolution persists when switching to another. The multi-resolution digital state is obtained by augmenting the reference digital state with the resolution-specific digital states. Conditioning the PGM's update and QoI factors on the digital action yields a unified and consistent Bayesian framework for multi-resolution state estimation and prediction.

The application of the multi-resolution digital twin to a multiscale system shows the computational efficiency and the accuracy of the proposed method. The multiscale system used a low-resolution static model of a macroscale finite element representation and a high-resolution static model using a nested FE² homogenization scheme with mesoscale and microscale descriptions for degrading composite panel. Using an unscented Kalman filter for sequential state estimation, we examined three multi-resolution strategies: a purely coarse-resolution twin, a fully high-resolution twin, and a hybrid approach that intermittently activated the high-fidelity model. The results demonstrate that sparse high-resolution updates improve state estimation accuracy without incurring the prohibitive cost of continuous multiscale analysis. The hybrid strategy effectively tracked both gradual and abrupt degradation events, preserved predictive accuracy close to that of the fully high-resolution solution, and significantly reduced computational effort. These findings illustrate that multi-resolution strategies can offer a favorable balance between fidelity and efficiency in practical monitoring scenarios.

A natural extension of this work is using adaptive resolution policies that rely on uncertainty quantification, value-of-information metrics, or risk-based criteria, moving beyond the predefined digital action schedule adopted in the present study. Another direction, within the multiscale multi-resolution context, is the joint inference of degradation constitutive law parameters alongside degradation states, enabling the digital twin to learn evolving material behavior; thereby enabling predictive capability to aid in decision-making regarding preventive maintenance as damage progresses. Lastly, the proposed multi-resolution architecture is also general enough to incorporate reduced-order models, multifidelity surrogates, or data-driven operators as additional resolutions, expanding the model resolution space available to the digital twin and is a promising direction for future investigation.

A. Mori–Tanaka formulation of the microscale

The Mori–Tanaka method provides an analytical estimate of the effective elasticity tensor of a composite comprising a matrix phase and one or more families of inclusion phases [35]. Since this formulation is time-independent, the timestep index t is omitted for brevity in the expressions that follow. For an inclusion family i with stiffness tensor \mathbf{C}_i embedded in a matrix of stiffness \mathbf{C}_{mat} , the strain concentration tensor is given by

$$\mathbf{B}_i = [\mathbf{I} + \mathbf{S}_i : \mathbf{C}_{\text{mat}}^{-1} : (\mathbf{C}_i - \mathbf{C}_{\text{mat}})]^{-1}, \quad (22)$$

where \mathbf{S}_i is the Eshelby tensor of an ellipsoidal inclusion of family i embedded in an infinite matrix.

The effective stiffness estimate is then

$$\mathbf{C}_{\text{MT}} = \mathbf{C}_{\text{mat}} + \sum_i v_i (\mathbf{C}_i - \mathbf{C}_{\text{mat}}) : \mathbf{B}_i : \left(v_{\text{mat}} \mathbf{I} + \sum_j v_j \mathbf{B}_j \right)^{-1}, \quad (23)$$

where v_i is the volume fraction of inclusion family i , and $v_{\text{mat}} = 1 - \sum_i v_i$ is the matrix volume fraction.

Prolate ellipsoidal inclusions for aligned fibers

The carbon fibers are modeled as aligned, highly elongated inclusions. Their geometry is represented by a prolate spheroid with semi-axes $\alpha_1 \gg \alpha_2 = \alpha_3$, aligned with the fiber direction. For such inclusions, the Eshelby tensor \mathbf{S}_{fib}

admits a closed-form expression in terms of the aspect ratio $\xi = \alpha_1/\alpha_2$. In the local coordinate system whose first axis is aligned with the fiber direction, the nonzero components take the form

$$S_{\text{fib}}^{11} = \frac{1}{\xi^2 - 1} \left(1 - \frac{\xi}{2\sqrt{\xi^2 - 1}} \ln \frac{\xi + \sqrt{\xi^2 - 1}}{\xi - \sqrt{\xi^2 - 1}} \right),$$

with the transverse components $S_{\text{fib}}^{22} = S_{\text{fib}}^{33}$ constructed to satisfy the Eshelby trace condition. The remaining components follow from symmetry. This geometric anisotropy enables Mori–Tanaka to capture the strong longitudinal stiffness characteristic of unidirectional fiber-reinforced composites.

Spherical inclusions for voids

Void growth is represented by a second inclusion family consisting of nearly spherical cavities. For a spherical inclusion, the Eshelby tensor simplifies and becomes isotropic:

$$S_{\text{void}} = \frac{3}{3K_{\text{mat}} + 4G_{\text{mat}}} \left(K_{\text{mat}} \mathbf{I}^{\text{vol}} + \frac{4}{3} G_{\text{mat}} \mathbf{I}^{\text{dev}} \right),$$

where K_{mat} and G_{mat} are the bulk and shear moduli of the matrix, and \mathbf{I}^{vol} and \mathbf{I}^{dev} are the volumetric and deviatoric projection tensors. The spherical geometry leads to an isotropic softening contribution, consistent with experimental observations of porosity-induced degradation.

Multiple inclusion families

The composite in the present application contains two inclusion families:

- 1) elongated carbon fibers, modeled as prolate spheroids with a highly anisotropic Eshelby tensor \mathbf{S}_{fib} ;
- 2) voids, modeled as spherical inclusions with isotropic \mathbf{S}_{void} .

The Mori–Tanaka estimate naturally combines these families through the weighted concentration tensors $\nu_i \mathbf{B}_i$ and the denominator $(\nu_{\text{mat}} \mathbf{I} + \sum_i \nu_i \mathbf{B}_i)^{-1}$. This multi-family formulation enables the model to represent anisotropic stiffening due to aligned fibers along with isotropic softening due to void growth.

Acknowledgment

The authors acknowledge support from the AFOSR grant FA9550-24-1-0327 under the Multidisciplinary University Research Initiatives (MURI), the AFOSR grant FA9550-23-1-0678, the NSF FDT-Biotech award 2436499, and the Texas Institute for Electronics (TIE) DARPA #HR00112430347.

References

- [1] National Academies of Sciences, Engineering, and Medicine, *Foundational research gaps and future directions for digital twins*, The National Academies Press, Washington, DC, 2023.
- [2] Ferrari, A., and Willcox, K., “Digital twins in mechanical and aerospace engineering,” *Nature Computational Science*, Vol. 4, No. 3, 2024, pp. 178–183.
- [3] Rasheed, A., San, O., and Kvamsdal, T., “Digital twin: Values, challenges and enablers from a modeling perspective,” *IEEE access*, Vol. 8, 2020, pp. 21980–22012.
- [4] Thelen, A., Zhang, X., Fink, O., Lu, Y., Ghosh, S., Youn, B. D., Todd, M. D., Mahadevan, S., Hu, C., and Hu, Z., “A comprehensive review of digital twin—part 1: modeling and twinning enabling technologies,” *Structural and Multidisciplinary Optimization*, Vol. 65, No. 12, 2022, p. 354.
- [5] Grieves, M., and Vickers, J., *Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems*, Springer International Publishing, Cham, 2017.
- [6] Glaessgen, E., and Stargel, D., “The digital twin paradigm for future NASA and US Air Force vehicles,” *53rd AIAA/ASME/ASCE/AHS/ASC structures, structural dynamics and materials conference 20th AIAA/ASME/AHS adaptive structures conference 14th AIAA*, 2012, p. 1818.

[7] Tuegel, E. J., Ingraffea, A. R., Eason, T. G., and Spottswood, S. M., “Reengineering aircraft structural life prediction using a digital twin,” *International Journal of Aerospace Engineering*, Vol. 2011, No. 1, 2011, p. 154798.

[8] Kapteyn, M. G., Pretorius, J. V., and Willcox, K. E., “A probabilistic graphical model foundation for enabling predictive digital twins at scale,” *Nature Computational Science*, Vol. 1, No. 5, 2021, pp. 337–347.

[9] McClellan, A., Lorenzetti, J., Pavone, M., and Farhat, C., “A physics-based digital twin for model predictive control of autonomous unmanned aerial vehicle landing,” *Philosophical Transactions of the Royal Society A*, Vol. 380, No. 2229, 2022, p. 20210204.

[10] Karve, P. M., Guo, Y., Kapusuzoglu, B., Mahadevan, S., and Haile, M. A., “Digital twin approach for damage-tolerant mission planning under uncertainty,” *Engineering Fracture Mechanics*, Vol. 225, 2020, p. 106766.

[11] Jiang, F., Ma, L., Broyd, T., and Chen, K., “Digital twin and its implementations in the civil engineering sector,” *Automation in Construction*, Vol. 130, 2021, p. 103838.

[12] Ye, C., Butler, L., Bartek, C., Iangurazov, M., Lu, Q., Gregory, A., Girolami, M., and Middleton, C., “A digital twin of bridges for structural health monitoring,” *12th International Workshop on Structural Health Monitoring*, Stanford University, 2019.

[13] Zhang, Z. J., Bakar, A., Humphry, A., Javid, F., Nadeau, P., Ebrahimi, M., Butscher, A., Tessier, A., Rodriguez, J., and Farhat, C., “Probabilistic learning from real-world observations of systems with unknown inputs for model-form UQ and digital twinning,” *Computer Methods in Applied Mechanics and Engineering*, Vol. 440, 2025, p. 117863.

[14] Laubenbacher, R., Mehrad, B., Shmulevich, I., and Trayanova, N., “Digital twins in medicine,” *Nature Computational Science*, Vol. 4, No. 3, 2024, pp. 184–191.

[15] Chaudhuri, A., Pash, G., Hormuth, D. A., Lorenzo, G., Kapteyn, M., Wu, C., Lima, E. A. B. F., Yankeelov, T. E., and Willcox, K., “Predictive digital twin for optimizing patient-specific radiotherapy regimens under uncertainty in high-grade gliomas,” *Frontiers in Artificial Intelligence*, Vol. 6, 2023, p. 1222612.

[16] Sel, K., Hawkins-Daarud, A., Chaudhuri, A., Osman, D., Bahai, A., Paydarfar, D., Willcox, K., Chung, C., and Jafari, R., “Survey and perspective on verification, validation, and uncertainty quantification of digital twins for precision medicine,” *npj Digital Medicine*, Vol. 8, No. 1, 2025, p. 40.

[17] Peherstorfer, B., Willcox, K., and Gunzburger, M., “Survey of multifidelity methods in uncertainty propagation, inference, and optimization,” *Siam Review*, Vol. 60, No. 3, 2018, pp. 550–591.

[18] Kapteyn, M., Knezevic, D., Huynh, D., Tran, M., and Willcox, K., “Data-driven physics-based digital twins via a library of component-based reduced-order models,” *International Journal for Numerical Methods in Engineering*, Vol. 123, No. 13, 2022, pp. 2986–3003.

[19] Lai, X., Yang, L., He, X., Pang, Y., Song, X., and Sun, W., “Digital twin-based structural health monitoring by combining measurement and computational data: An aircraft wing example,” *Journal of Manufacturing Systems*, Vol. 69, 2023, pp. 76–90.

[20] Zhang, Y., Hao, R., Niiranen, J., Yang, Y., Brühwiler, E., Su, D., and Nagayama, T., “Computational engine for finite element digital twins of structural dynamics via motion data,” *Engineering Structures*, Vol. 316, 2024, p. 118630.

[21] Corbetta, M., Sbarufatti, C., Giglio, M., and Todd, M. D., “Optimization of nonlinear, non-Gaussian Bayesian filtering for diagnosis and prognosis of monotonic degradation processes,” *Mechanical Systems and Signal Processing*, Vol. 104, 2018, pp. 305–322.

[22] Cai, B., Huang, L., and Xie, M., “Bayesian networks in fault diagnosis,” *IEEE Transactions on Industrial Informatics*, Vol. 13, No. 5, 2017, pp. 2227–2240.

[23] Bartram, G., and Mahadevan, S., “Probabilistic prognosis with dynamic bayesian networks,” *International Journal of Prognostics and Health Management*, Vol. 6, No. 4, 2015.

[24] Ghahramani, Z., “Learning dynamic Bayesian networks,” *International School on Neural Networks, Initiated by IIASS and EMFCSC*, 1997, pp. 168–197.

[25] Murphy, K. P., “Dynamic Bayesian Networks: Representation, Inference and Learning,” Ph.D. thesis, University of California, Berkeley, 2002.

[26] Li, C., Mahadevan, S., Ling, Y., Choze, S., and Wang, L., “Dynamic Bayesian network for aircraft wing health monitoring digital twin,” *AIAA Journal*, Vol. 55, No. 3, 2017, pp. 930–941.

- [27] Geers, M. G., Kouznetsova, V. G., Matouš, K., and Yvonnet, J., “Homogenization methods and multiscale modeling: nonlinear problems,” *Encyclopedia of Computational Mechanics Second Edition*, 2017, pp. 1–34.
- [28] Matouš, K., Geers, M. G., Kouznetsova, V. G., and Gillman, A., “A review of predictive nonlinear theories for multiscale modeling of heterogeneous materials,” *Journal of Computational Physics*, Vol. 330, 2017, pp. 192–220.
- [29] Julier, S. J., and Uhlmann, J. K., “Unscented filtering and nonlinear estimation,” *Proceedings of the IEEE*, Vol. 92, No. 3, 2004, pp. 401–422.
- [30] Miehe, C., and Koch, A., “Computational micro-to-macro transitions of discretized microstructures undergoing small strains,” *Archive of Applied Mechanics*, Vol. 72, No. 4, 2002, pp. 300–317.
- [31] Feyel, F., “A multilevel finite element method (FE2) to describe the response of highly non-linear structures using generalized continua,” *Computer Methods in Applied Mechanics and Engineering*, Vol. 192, No. 28–30, 2003, pp. 3233–3244.
- [32] Mori, T., and Tanaka, K., “Average stress in matrix and average elastic energy of materials with misfitting inclusions,” *Acta Metallurgica*, Vol. 21, No. 5, 1973, pp. 571–574.
- [33] Yang, J. N., and Liu, M., “Residual strength degradation model and theory of periodic proof tests for graphite/epoxy laminates,” *Journal of Composite Materials*, Vol. 11, No. 2, 1977, pp. 176–203.
- [34] Pirondi, A., and Moroni, F., “Improvement of a cohesive zone model for fatigue delamination rate simulation,” *Materials*, Vol. 12, No. 1, 2019, p. 181.
- [35] Mura, T., *Micromechanics of Defects in Solids*, reprint ed., Springer Science & Business Media, 2013.