


Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Comput. Methods Appl. Mech. Engrg.

journal homepage: www.elsevier.com/locate/cma

A digital twin framework for civil engineering structures

Matteo Torzoni ^{a,*}, Marco Tezzele ^b, Stefano Mariani ^a, Andrea Manzoni ^c,
Karen E. Willcox ^b^a Dipartimento di Ingegneria Civile e Ambientale, Politecnico di Milano, Piazza L. da Vinci 32, Milan, 20133, Italy^b Oden Institute for Computational Engineering and Sciences, University of Texas at Austin, Austin, 78712, TX, United States^c MOX, Dipartimento di Matematica, Politecnico di Milano, Piazza L. da Vinci 32, Milan, 20133, Italy

ARTICLE INFO

Keywords:

Digital twins
Predictive maintenance
Bayesian networks
Deep learning
Structural health monitoring
Model order reduction

ABSTRACT

The digital twin concept represents an appealing opportunity to advance condition-based and predictive maintenance paradigms for civil engineering systems, thus allowing reduced lifecycle costs, increased system safety, and increased system availability. This work proposes a predictive digital twin approach to the health monitoring, maintenance, and management planning of civil engineering structures. The asset-twin coupled dynamical system is encoded employing a probabilistic graphical model, which allows all relevant sources of uncertainty to be taken into account. In particular, the time-repeating observations-to-decisions flow is modeled using a dynamic Bayesian network. Real-time structural health diagnostics are provided by assimilating sensed data with deep learning models. The digital twin state is continually updated in a sequential Bayesian inference fashion. This is then exploited to inform the optimal planning of maintenance and management actions within a dynamic decision-making framework. A preliminary offline phase involves the population of training datasets through a reduced-order numerical model and the computation of a health-dependent control policy. The strategy is assessed on two synthetic case studies, involving a cantilever beam and a railway bridge, demonstrating the dynamic decision-making capabilities of health-aware digital twins.

1. Introduction

The optimal management of deteriorating structural systems is an important challenge in modern engineering. In particular, the failure or non-optimized maintenance planning of civil structures may entail high safety, economic, and social costs. Within this context, enabling a digital twin (DT) perspective for structural systems that are critical for either safety or operative reasons, is crucial to allow for condition-based or predictive maintenance practices, in place of customarily employed time-based ones. Indeed, having an up-to-date digital replica of the physical asset of interest can yield several benefits spanning its entire lifecycle, including performance and health monitoring, as well as maintenance, inspection, and management planning [1].

The DT concept [2–6] has been recently applied to several fields for operational monitoring, control, and decision support, including structural health monitoring (SHM) and predictive maintenance [7,8], additive manufacturing [9], smart cities [10], urban sustainability [11], and railway systems management [12]. It allows for a personalized characterization of a physical asset, in the form of computational models and parameters of interest, that evolves over time and is kept synchronized with its physical counterpart by means of data-collecting devices. Within a civil SHM framework, such a twinning perspective can be enabled by

* Corresponding author.

E-mail addresses: matteo.torzoni@polimi.it (M. Torzoni), marco.tezzele@austin.utexas.edu (M. Tezzele), stefano.mariani@polimi.it (S. Mariani), andreal.manzoni@polimi.it (A. Manzoni), kwillcox@oden.utexas.edu (K.E. Willcox).<https://doi.org/10.1016/j.cma.2023.116584>

Received 2 August 2023; Received in revised form 30 October 2023; Accepted 31 October 2023

Available online 8 November 2023

0045-7825/© 2023 The Author(s).

Published by Elsevier B.V. This is an open access article under the CC BY license

<http://creativecommons.org/licenses/by/4.0/>.

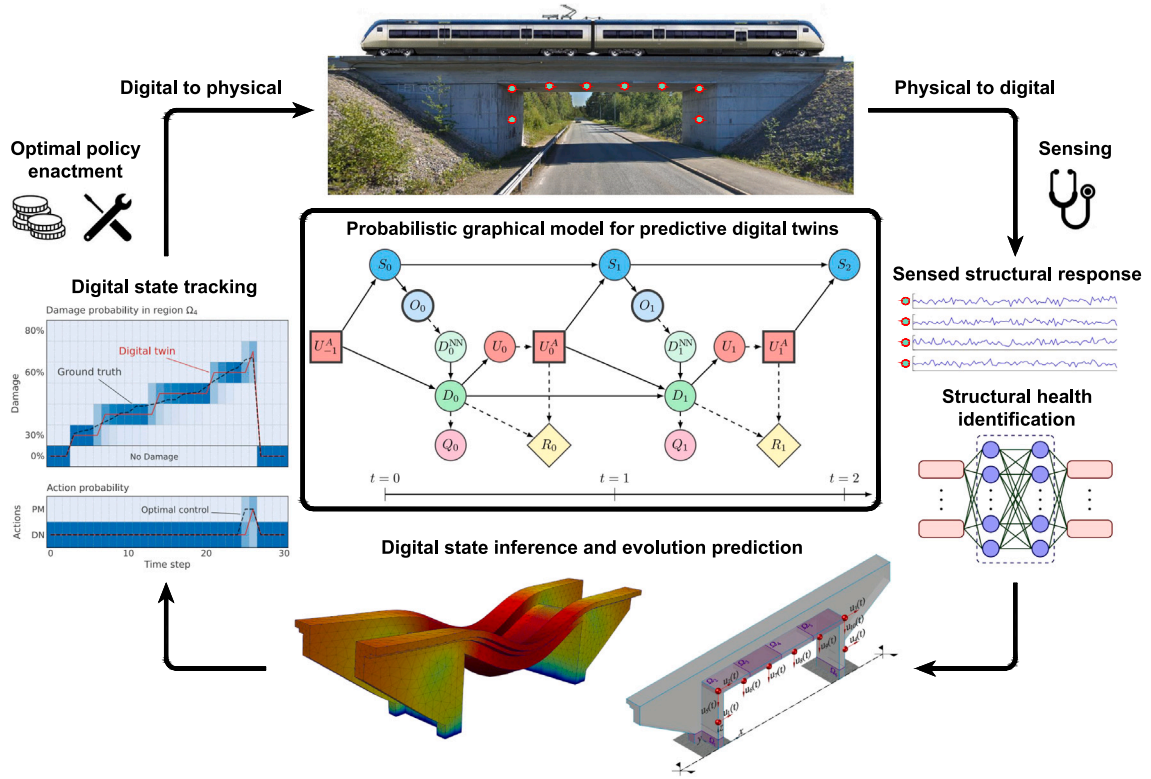


Fig. 1. Predictive digital twin framework for civil engineering structures: graphical abstraction of the end-to-end information flow enabled by the probabilistic graphical model.

the assimilation of data through data-driven structural health diagnostics (from physical to digital), possibly accommodating the quantification and propagation of relevant uncertainties related to, e.g., measurement noise, modeling assumptions, environmental and operational variabilities [13–17]. The resulting updated digital state should then enable prediction of the physical system evolution, as well as inform optimal planning of maintenance and management actions (from digital to physical).

In this work, we propose a DT framework for civil engineering structures. The overall computational strategy is based upon a probabilistic graphical model (PGM) inspired by the foundational model proposed in [18], which provides a general framework to carry out data assimilation, state estimation, prediction, planning, and learning. Formally, such a PGM is a dynamic Bayesian network with the addition of decision nodes, i.e., a dynamic decision network [19,20]. This is employed to encode the end-to-end information flow, from physical to digital through assimilation and inference, and back to the physical asset in the form of informed control actions. A graphical abstraction of the proposed DT strategy is depicted in Fig. 1. The figure shows a *physical-to-digital* information flow and a *digital-to-physical* information flow. These bi-directional information flows repeat indefinitely over time. In particular, we have:

- *From physical to digital.* Structural response data are gathered from the physical system and assimilated with deep learning (DL) models, see e.g., [21,22], to estimate the current structural health in terms of presence, location, and severity of structural damage. To solve this inverse problem, we refer to vibration-based SHM techniques, see e.g., [23–26], which exploit the aforementioned collected data, such as displacement or acceleration time histories. This first estimate of the digital state is then employed to estimate an updated digital state, according to control-dependent transition dynamics models describing how the structural health is expected to evolve.
- *From digital to physical.* The updated digital state is exploited to predict the future evolution of the physical system and the associated uncertainty, thereby enabling predictive decision-making about maintenance and management actions feeding back to the physical system.
- *Offline learning phase.* The DT setup considered in this work takes advantage of a preliminary offline learning phase. This phase involves the training of the DL models underlying the structural health identification, and learning the control policy to be applied at each time step of the online phase. The DL models are trained in a supervised fashion, with labeled data pertaining to specific damage conditions generated by exploiting physics-based numerical models. To efficiently assemble a training dataset representative of potential damage and operational conditions the structure might undergo during its lifetime, we exploit a reduced-order modeling strategy for parametrized systems relying on the reduced basis method [27]. The health-dependent

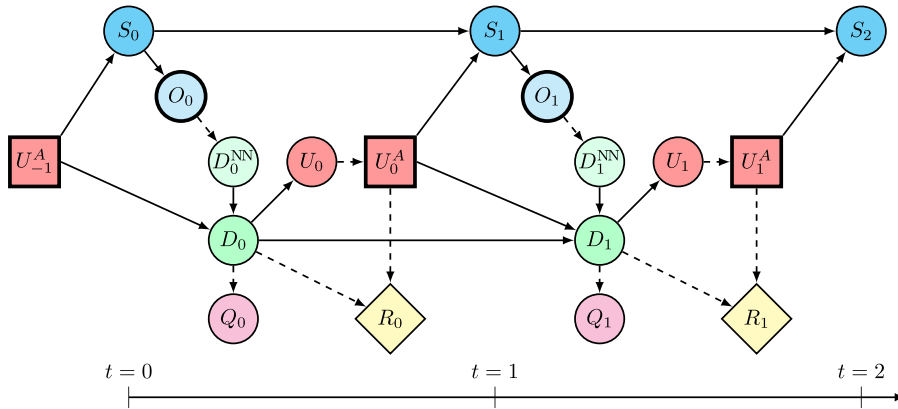


Fig. 2. Dynamic decision network encoding the asset-twin coupled dynamical system. Circle nodes denote random variables, square nodes denote actions, and diamond nodes denote the objective function. Bold outlines denote observed quantities, while thin outlines denote estimated quantities. Directed solid edges represent the variables' dependencies encoded via conditional probability distributions, while directed dashed edges represent the variables' dependencies encoded via deterministic functions.

control policy is also computed offline, by maximizing the expected future rewards for the planning problem induced by the PGM.

The elements of novelty that characterize this work are the following: (i) the adaptation of the PGM digital twinning framework to the health monitoring, maintenance, and management planning of civil engineering structures; (ii) the assimilation of vibration response data is carried out by exploiting DL models, which allow automated selection and extraction of optimized damage-sensitive features and real-time assessment of the structural state. This work shows how to incorporate in the DT framework high-dimensional multivariate time series describing the sensor measurements, while tracking the associated uncertainties. The proposed computational framework is made available in the public repository `digital-twin-SHM` [28]. The code implements the PGM framework as a dynamic decision network. This enables us to easily specify the graph topology from a few time slices, and then unroll it for any number of time steps in the future.

The remainder of this paper is organized as follows. In Section 2, we describe the proposed DT framework. In Section 3, the computational procedure is assessed on two test cases, respectively related to a cantilever beam and a railway bridge. Conclusions and future developments are drawn in Section 4.

2. Predictive digital twins using physics-based models and machine learning

In this section, we describe the methodology characterizing our DT framework in terms of the PGM encoding the asset-twin coupled dynamical system in Section 2.1; the population of training datasets exploiting physics-based numerical models in Section 2.2; and the DL models underlying the structural health identification in Section 2.3.

2.1. Probabilistic graphical model for predictive digital twins

The digital twin assimilates vibration recordings shaped as multivariate time series $\mathbf{U}(\boldsymbol{\mu}) = [\mathbf{u}_1(\boldsymbol{\mu}), \dots, \mathbf{u}_{N_u}(\boldsymbol{\mu})] \in \mathbb{R}^{L \times N_u}$, consisting of N_u time series made of L sensor measurements equally spaced in time, for instance in terms of accelerations or displacements. The vector $\boldsymbol{\mu} \in \mathbb{R}^{N_{\text{par}}}$ comprises the parameters representing the operational, damage, and (possibly) environmental conditions. Each recording refers to a time interval $(0, T)$, within which measurements are recorded with a sampling rate f_s . For the problem settings we consider, the interval $(0, T)$ is short enough for the operational, environmental, and damage conditions to be considered time-invariant, yet long enough not to compromise the identification of the structural behavior.

The PGM that defines the elements comprising the asset-twin coupled dynamical system, and mathematically describes the relevant interactions via observed data and control inputs, is the dynamic decision network sketched in Fig. 2. Circle nodes in the graph denote random variables at discrete times, square nodes denote actions, and diamond nodes denote the objective function. Bold outlines denote observed quantities, while thin outlines denote estimated quantities. The directed acyclic structure of the PGM encodes the assumed conditional dependencies. Edges in the graph represent dependencies between random variables. Solid edges represent the variables' dependencies encoded via conditional probability distributions, while dashed edges represent the variables' dependencies encoded via deterministic functions.

We consider a non-dimensional time discretization, and denote discrete time steps by t . The physical time duration between successive time steps may vary depending on the application, and is governed by the update frequency of the DT via data assimilation, so that the DT is updated once per time step. Thanks to the modeled conditional dependencies between random variables, the graph topology is specified from the first two time steps, and can then be unrolled for any number of time steps.

The physical state $S_t \sim p(s_t)$, with s_t denoting the realization of the random variable S_t at time t , encapsulates the variability in the health state of the asset, which is usually only partially observable. The probability distribution encoding the relative likelihood that $S_t = s_t$, for any possible s_t , is denoted with $p(s_t)$. The digital state $D_t \sim p(d_t)$ is characterized by those parameters employed to capture the variability of the physical asset by means of the computational models comprising the DT. In our framework, the digital state is given as a vector of length two, describing the presence/location and magnitude of damage in the asset. The *physical-to-digital* information flow is governed by the observed data $O_t = o_t$, which are assimilated by the DT to update the digital state. The assimilation is carried out using the DL models described in Section 2.3, providing a first estimate of the digital state $D_t^{\text{NN}} \sim p(d_t^{\text{NN}})$. This estimation is then used in a Bayesian inference formulation, together with the prior belief D_{t-1} from the previous time step, to estimate an updated digital state D_t according to a control-dependent transition dynamics model describing how the digital state is expected to evolve. The updated digital state can thus be exploited to compute quantities of interest $Q_t \sim p(q_t)$, such as modal quantities or other response features, through the computational models comprising the DT. For instance, quantities of interest can be useful to perform posterior predictive checks on the tracking capabilities of the DT to assess how it matches the reality, by comparing sensor measurements with the corresponding posterior estimates. However, we point out that this capability is not exploited in the present work, and that the Q_t node is kept in the graph in agreement with the foundational model proposed in [18]. Nevertheless, the updated digital state D_t is eventually exploited to inform the *digital-to-physical* information flow, in the form of control inputs; in Fig. 2, $U_t \sim p(u_t)$ and $U_t^A = u_t^A$ denote the belief about what action to take and the control input effectively enacted on the asset, respectively. At each time step, U_t is estimated according to a health-dependent control policy, that maps the belief over the digital state onto the control actions feeding back to the physical asset. Finally, the reward $R_t \sim p(r_t)$ quantifies the performance of the asset for the time step and can be equivalently perceived as a negative cost to be maximized.

The key assumptions behind our PGM are that the physical state is only observable indirectly via the sensed structural response, and the physical and digital states evolve according to a Markovian process. This implies that the conditional probabilities associated with the random variables at one time step depend only on the random variables at the previous time step, and are independent of all past states. The resulting graph topology encodes a conditional independence structure that allows us to conveniently cast the asset tracking within a sequential Bayesian inference framework. Indeed, by exploiting conditional independence and Bayes rule, the joint distributions over variables can be factorized up to the current time step t_c , as follows:

$$p(D_0^{\text{NN}}, \dots, D_{t_c}^{\text{NN}}, D_0, \dots, D_{t_c}, Q_0, \dots, Q_{t_c}, R_0, \dots, R_{t_c}, U_0, \dots, U_{t_c} | o_0, \dots, o_{t_c}, u_0^A, \dots, u_{t_c}^A) \\ \propto \prod_{t=0}^{t_c} \left[\phi_t^{\text{data}} \phi_t^{\text{history}} \phi_t^{\text{NN}} \phi_t^{\text{QoI}} \phi_t^{\text{control}} \phi_t^{\text{reward}} \right], \quad (1)$$

with factors:

$$\phi_t^{\text{data}} = p(O_t = o_t | D_t^{\text{NN}}), \quad (2)$$

$$\phi_t^{\text{history}} = p(D_t | D_{t-1}, U_{t-1}^A = u_{t-1}^A), \quad (3)$$

$$\phi_t^{\text{NN}} = p(D_t | D_t^{\text{NN}}), \quad (4)$$

$$\phi_t^{\text{QoI}} = p(Q_t | D_t), \quad (5)$$

$$\phi_t^{\text{control}} = p(U_t | D_t), \quad (6)$$

$$\phi_t^{\text{reward}} = p(R_t | D_t, U_t^A = u_t^A). \quad (7)$$

Herein, ϕ_t^{data} encodes the assimilation of observed data through the DL models underlying the identification of the structural health. ϕ_t^{history} and ϕ_t^{NN} factorize the belief about the digital state D_t , conditioned on the digital state at the previous time step D_{t-1} , the last enacted control input $U_{t-1}^A = u_{t-1}^A$, and the data assimilation outcome D_t^{NN} . In our PGM, the spaces of the digital states and control inputs are discrete, thus the relevant causal relationships are modeled by means of conditional probability tables (CPTs). In particular, ϕ_t^{history} plays the role of a predictor forward in time, parameterized by means of a control-dependent CPT describing how the digital state is expected to evolve. Such a CPT should embody any a priori knowledge that the DT designer has with respect to the asset and the relevant operational conditions. ϕ_t^{history} can be estimated offline from historical data, see e.g., [29,30], or learned online from experience. On the other hand, ϕ_t^{NN} updates the digital state estimate to account for data assimilation. This is encoded by means of a CPT mapping the estimate D_t^{NN} provided by the DL models, onto a belief about D_t . Such a CPT is a confusion matrix measuring the offline (expected) performance of the DL models in correctly identifying the digital state among all the possible outcomes of D_t . ϕ_t^{QoI} and ϕ_t^{reward} respectively encapsulate the evaluation of the computational models comprising the DT to estimate quantities of interest, and the computation of the reward function quantifying the performance of the asset. Finally, the control factor ϕ_t^{control} is encoded by means of a health-dependent control policy $\pi(D_t)$ computed as described in the following. Since the spaces of the unobserved variables are discrete, we can propagate and update the relative belief exactly with a single pass of the sum-product message-passing algorithm [19].

The control policy $\pi(D_t)$ is computed offline under the simplifying assumption of sufficient sensing capability to provide an accurate estimate of the structural health, allowing us to decouple the sensing and control problems. This involves solving the planning problem induced by the expected evolution of the structural health, maximizing the expected reward over the planning horizon. Considering an infinite planning horizon, this can be stated as the optimization problem:

$$\pi(D_t) = \arg \max_{\pi} \sum_{t=0}^{+\infty} \gamma^t \mathbb{E}[R_t], \quad (8)$$

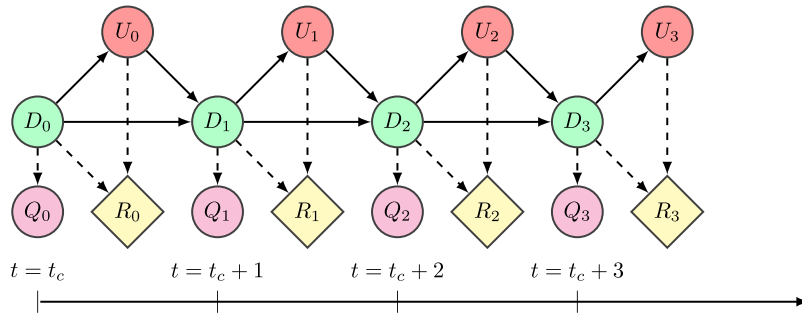


Fig. 3. Dynamic decision network employed to predict the future evolution of the digital state and the associated uncertainty. Circle nodes denote random variables, and diamond nodes denote the objective function. Directed solid edges represent the variables' dependencies encoded via conditional probability distributions, while directed dashed edges represent the dependencies encoded via deterministic functions.

where $\gamma \in [0, 1]$ is the discount factor. Here, this is solved using the dynamic-programming value iteration algorithm [31]. The reward function to be optimized is chosen as:

$$R_t(U_t, D_t) = R_t^{\text{control}}(U_t) + \alpha R_t^{\text{health}}(D_t). \quad (9)$$

Herein, $R_t^{\text{control}}(U_t)$ and $R_t^{\text{health}}(D_t)$ quantify the rewards relative to control inputs and health state, respectively, and $\alpha \in \mathbb{R}$ is a weighting factor, useful to tune the trade-off between risk-averse and risk-seeking behavior. After learning $\pi(D_t)$, U_t^A is selected as the best point estimate of U_t .

Starting from the updated digital state D_{t_c} at the current time step t_c , future prediction is achieved by unrolling until a prediction time t_p the portion of PGM relative to D_t , Q_t , R_t , and U_t (see Fig. 3). All other nodes are removed from the prediction graph, as neither data assimilation nor actions are performed on the asset while forecasting its evolution. The factorization in Eq. (1) can be extended over the prediction horizon as:

$$p(D_0^{\text{NN}}, \dots, D_{t_c}^{\text{NN}}, D_0, \dots, D_{t_p}, Q_0, \dots, Q_{t_p}, R_0, \dots, R_{t_p}, U_0, \dots, U_{t_p} | o_0, \dots, o_{t_c}, u_0^A, \dots, u_{t_c}^A) \propto \prod_{t=0}^{t_p} [\phi_t^{\text{history}} \phi_t^{\text{QoI}} \phi_t^{\text{control}} \phi_t^{\text{reward}}] \prod_{t=0}^{t_c} [\phi_t^{\text{data}} \phi_t^{\text{NN}}]. \quad (10)$$

The algorithmic description of the online phase of the proposed digital twinning framework is reported in Algorithm 1. The operations repeat each time new observational data are provided. Note that the considered PGM digital twinning framework is general, and can easily be adapted to deal with physical assets other than civil engineering structures by reorganizing the topology of the graph, if necessary.

Algorithm 1 Online phase – algorithmic description

Input: observational data $O_t = o_t$

- 1: assimilate o_t with the DL models to provide $D_t^{\text{NN}} = d_t^{\text{NN}}$. $\triangleright (O_t) \rightarrow (D_t^{\text{NN}})$
 - 2: infer D_t and U_t by updating d_{t-1} , given u_{t-1}^A , d_{t-1}^{NN} , and the CPTs encoding ϕ_t^{history} , ϕ_t^{NN} and ϕ_t^{control} . $\triangleright (D_{t-1}, D_{t-1}^{\text{NN}}, U_{t-1}^A) \rightarrow (D_t, U_t)$
 - 3: infer the future evolution of D_t and U_t , given the updated d_t , and the CPTs encoding ϕ_t^{history} and ϕ_t^{control} . $\triangleright (D_t) \rightarrow (D_{t_p}, U_{t_p})$
 - 4: select $U_t^A = u_t^A$ as the best point estimate of $U_t = u_t$. $\triangleright (U_t) \rightarrow (U_t^A)$
 - 5: **return** control input to be enacted u_t^A , expected evolution of D_t and U_t .
-

2.2. Numerical models for simulation-based damage identification

As anticipated in the previous section, the assimilation of structural response data to identify the structural state is carried out through DL models. A simulation-based strategy is exploited to train the DL models on the basis of vibration responses. The training data are numerically generated by simulating physics-based models so that the effect of damage on the structural response can be systematically reproduced [32]. In particular, the structure to be monitored is modeled as a linear-elastic continuum, discretized in space through finite elements. Its dynamic response to the applied loadings, under the assumption of linearized kinematics, is described by the following semi-discretized form of the elasto-dynamic problem:

$$\begin{cases} \mathbf{M}\ddot{\mathbf{x}}(t) + \mathbf{C}(\boldsymbol{\mu})\dot{\mathbf{x}}(t) + \mathbf{K}(\boldsymbol{\mu})\mathbf{x}(t) = \mathbf{f}(t, \boldsymbol{\mu}), & t \in (0, T) \\ \mathbf{x}(0) = \mathbf{x}_0, \\ \dot{\mathbf{x}}(0) = \dot{\mathbf{x}}_0, \end{cases} \quad (11)$$

which is referred to as the full-order model (FOM). Here $t \in (0, T)$ denotes time; $\mathbf{x}(t), \dot{\mathbf{x}}(t), \ddot{\mathbf{x}}(t) \in \mathbb{R}^{N_{FE}}$ are the vectors of nodal displacements, velocities and accelerations, respectively; N_{FE} is the number of degrees of freedom (dofs); $\mathbf{M} \in \mathbb{R}^{N_{FE} \times N_{FE}}$ is the mass matrix; $\mathbf{C}(\boldsymbol{\mu}) \in \mathbb{R}^{N_{FE} \times N_{FE}}$ is the damping matrix, assembled according to the Rayleigh's model; $\mathbf{K}(\boldsymbol{\mu}) \in \mathbb{R}^{N_{FE} \times N_{FE}}$ is the stiffness matrix; $\mathbf{f}(t, \boldsymbol{\mu}) \in \mathbb{R}^{N_{FE}}$ is the vector of nodal forces induced by the external loadings; and \mathbf{x}_0 and $\dot{\mathbf{x}}_0$ are the initial conditions (at $t = 0$), in terms of nodal displacements and velocities, respectively. The mass matrix \mathbf{M} is not a function of $\boldsymbol{\mu}$ because the mass properties of the structure are unaffected by the employed damage description or by the operational conditions. The solution of Problem (11) is advanced in time using the Newmark integration scheme (constant average acceleration method) [33], to provide $\mathbf{x}_l, \dot{\mathbf{x}}_l$ and $\ddot{\mathbf{x}}_l$, for $l = 1, \dots, L$, with \mathbf{x}_l being the vector of nodal displacements at time l .

With reference to civil structures, we focus on the early detection of damage patterns characterized by a small evolution rate, whose prompt identification can reduce lifecycle costs and increase the safety and availability of the structure. In this context, a localized reduction of the material stiffness stands as the simplest damage mechanism resulting from a time scale separation between damage growth and damage assessment, see e.g., [34–36]. Here, local stiffness reduction is obtained by parametrizing the stiffness matrix via two variables $y \in \mathbb{N}$ and $\delta \in \mathbb{R}$, included in the parameter vector $\boldsymbol{\mu}$, respectively describing the location and magnitude of the applied stiffness reduction, similarly to [37–39]. In particular, $y \in \{0, \dots, N_y\}$ labels the specific damage region, among a set of predefined N_y damage locations, where $y = 0$ identifies the damage-free baseline. The parameter $\delta \in \mathbb{R}$ describes the magnitude of the stiffness reduction taking place within the predesignated region associated with y .

As N_{FE} increases, the computational cost associated with the solution of the FOM for any sampled $\boldsymbol{\mu}$ also grows, and the generation of synthetic datasets becomes prohibitive. To address this challenge, a projection-based reduced-order model (ROM) is exploited in place of the FOM to speed up the offline dataset population phase, similarly to [38,39]. The ROM is obtained by a proper orthogonal decomposition (POD)-Galerkin reduced basis method [27,40–42]. This reduced-order modeling strategy is chosen because POD has been investigated and validated in the context of structural dynamics [43,44] and structural analysis [45,46], its appealing offline–online decoupling, and the availability of efficient criteria for the selection of POD basis functions. It is worth noting that alternative reduced-order modeling approaches can also be employed to alleviate the computational burden during the offline dataset generation. For instance, one could use spectral POD [47–49], or Grassmannian diffusion maps [50], as viable alternatives to the reduced basis method.

The ROM approximation to the solution of Problem (11) is obtained by linearly combining $N_{RB} \ll N_{FE}$ POD basis functions $\mathbf{w}_k \in \mathbb{R}^{N_{FE}}$, $k = 1, \dots, N_{RB}$, as $\mathbf{x}(t, \boldsymbol{\mu}) \approx \mathbf{W}\hat{\mathbf{x}}(t, \boldsymbol{\mu})$, where $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_{N_{RB}}] \in \mathbb{R}^{N_{FE} \times N_{RB}}$ is the basis matrix collecting the POD basis functions and $\hat{\mathbf{x}}(t, \boldsymbol{\mu}) \in \mathbb{R}^{N_{RB}}$ is the vector of unknown POD coefficients. By enforcing the orthogonality between the residual and the subspace spanned by the first N_{RB} POD modes through a Galerkin projection, the following N_{RB} -dimensional semi-discretized form is obtained:

$$\begin{cases} \mathbf{M}_r \ddot{\hat{\mathbf{x}}}(t) + \mathbf{C}_r(\boldsymbol{\mu}) \dot{\hat{\mathbf{x}}}(t) + \mathbf{K}_r(\boldsymbol{\mu}) \hat{\mathbf{x}}(t) = \mathbf{f}_r(t, \boldsymbol{\mu}), & t \in (0, T) \\ \hat{\mathbf{x}}(0) = \mathbf{W}^T \mathbf{x}_0, \\ \dot{\hat{\mathbf{x}}}(0) = \mathbf{W}^T \dot{\mathbf{x}}_0. \end{cases} \quad (12)$$

The solution of this reduced-order system is advanced in time using the same strategy employed for the FOM model, and then projected onto the original FOM space as $\mathbf{x}(t, \boldsymbol{\mu}) \approx \mathbf{W}\hat{\mathbf{x}}(t, \boldsymbol{\mu})$. Here, reduced matrices \mathbf{M}_r , \mathbf{C}_r , and \mathbf{K}_r , and the reduced vector \mathbf{f}_r play the same role as their high-fidelity counterparts, yet with dimension $N_{RB} \times N_{RB}$ instead of $N_{FE} \times N_{FE}$, according to the following relationships:

$$\begin{aligned} \mathbf{M}_r &\equiv \mathbf{W}^T \mathbf{M} \mathbf{W}, & \mathbf{C}_r(\boldsymbol{\mu}) &\equiv \mathbf{W}^T \mathbf{C}(\boldsymbol{\mu}) \mathbf{W}, \\ \mathbf{K}_r(\boldsymbol{\mu}) &\equiv \mathbf{W}^T \mathbf{K}(\boldsymbol{\mu}) \mathbf{W}, & \mathbf{f}_r(t, \boldsymbol{\mu}) &\equiv \mathbf{W}^T \mathbf{f}(t, \boldsymbol{\mu}). \end{aligned} \quad (13)$$

The basis matrix \mathbf{W} is obtained by POD, exploiting the so-called method of snapshots as follows. First, a snapshot matrix $\mathbf{S} = [\mathbf{x}_1, \dots, \mathbf{x}_{N_S}] \in \mathbb{R}^{N_{FE} \times N_S}$ is assembled from N_S solution snapshots, computed by integrating in time the FOM solution for different values of parameters $\boldsymbol{\mu}$. The computation of an optimal reduced basis is then carried out by factorizing \mathbf{S} through a singular value decomposition. We use a standard energy-based criterion to set the order N_{RB} of the approximation. For further details see, e.g., [21,23,27,42].

To populate the training dataset D , the parametric space of vector $\boldsymbol{\mu}$ is taken as uniformly distributed, and then sampled via the Latin hypercube rule. The number of samples is equal to the number I of instances collected in D as:

$$D = \{(\mathbf{U}_i, y_i, \delta_i)\}_{i=1}^I, \quad (14)$$

where the vibration recordings \mathbf{U}_i associated with the i -th sampling of $\boldsymbol{\mu}$, with $i = 1, \dots, I$, are labeled by the corresponding values of y_i and δ_i , and are obtained as follows. With reference to displacement recordings, nodal values in $(0, T)$ are first collected as $\mathbf{V}_i = [\mathbf{W}\hat{\mathbf{x}}_1, \dots, \mathbf{W}\hat{\mathbf{x}}_L]_i \in \mathbb{R}^{N_{FE} \times L}$ by solving Problem (12). The relevant vibration recordings \mathbf{U}_i are then obtained as:

$$\mathbf{U}_i = (\mathbf{T}\mathbf{V}_i)^T, \quad (15)$$

where $\mathbf{T} \in \mathbb{R}^{N_u \times N_{FE}}$ is a Boolean matrix whose (n, m) -th entry is equal to 1 only if the n -th sensor output coincides with the m -th dof. In order to mimic the measurement noise, each vibration recording in D is corrupted by adding an independent, identically distributed Gaussian noise, whose statistical properties depend on the target accuracy of the sensors. In the following, the index i will be dropped for ease of notation, unless necessary.

2.3. Data assimilation via artificial neural networks

The ϕ_i^{data} factor in our PGM encodes the assimilation of observed data through the DL models underlying the identification of the structural health. In this section, we describe the adopted DL models, the aspects related to their training, and how they are used to assimilate observational data to detect, locate, and quantify the presence of structural damage.

Every time new observational data \mathbf{U} are acquired, they are first processed with a classification model NN_{CL} to address damage detection/localization. Classification involves the prediction of an output class to categorize a given input. Here, the classes are those described through the y parameter. Whenever a damage is identified in the j -th region, $j = 1, \dots, N_y$, the observational data \mathbf{U} are further processed with regression models NN_{RG}^j , one for each damageable region, to quantify the associated amount of damage δ .

The aforementioned classification and regression tasks are addressed by means of DL models. The use of DL models for SHM purposes has the advantage of automating the feature engineering stage characterizing the pattern recognition paradigm for SHM [35,51]. Indeed, a DL model is trained to select and extract optimized damage-sensitive features from raw sensor recordings through an end-to-end learning process. Moreover, since the DL model is learned offline, the structural state can be next assessed in real-time regardless of considering continuous or discrete variables, which would be difficult to achieve with other optimization techniques, such as nonlinear programming, stochastic optimization, and metaheuristic methods.

The model NN_{CL} addresses the multi-class classification task underlying the damage detection/localization problem, namely $\text{NN}_{\text{CL}} : \mathbf{U} \rightarrow \mathbf{b} \in \mathbb{R}^{N_y+1}$. The target label \mathbf{b} categorizes one of the $N_y + 1$ predefined damage scenarios described through parameter y . In particular, \mathbf{b} is a one-hot encoding $\mathbf{b} = [b^0, \dots, b^{N_y}]^T$, with entries b^m equal to 1 if the target class y is m and 0 otherwise, with $m = 0, \dots, N_y$. This is needed because DL models cannot operate on nominal data directly. They require all input variables and output variables to be numeric. The one-hot encoding converts the nominal feature described by the y parameter into a multidimensional binary vector. The number of dimensions corresponds to the number of categories, and each category gets its dimension. Each category is encoded by mapping it to a vector in which the entry corresponding to the category's dimension is 1, and the rest are 0.

The estimated counterpart of \mathbf{b} is obtained as $\hat{\mathbf{b}} = \text{NN}_{\text{CL}}(\mathbf{U})$. By employing a Softmax activation function for the output layer of NN_{CL} , the entries of $\hat{\mathbf{b}} = (\hat{b}^0, \dots, \hat{b}^{N_y})^T \in \mathbb{R}^{N_y+1}$ are interpreted as the confidence levels \hat{b}^m by which \mathbf{U} is assigned to the m -th damage class, with $m = 0, \dots, N_y$. In particular, the Softmax activation function converts the real-valued vector $\mathbf{a} = (a^0, \dots, a^{N_y})^T \in \mathbb{R}^{N_y+1}$, provided by the output layer of NN_{CL} , into a discrete probability distribution as:

$$\hat{\mathbf{b}} = \text{Softmax}(\mathbf{a}), \quad \text{with} \quad \hat{b}^m(\mathbf{a}) = \frac{\exp(a^m)}{\sum_{k=0}^{N_y} \exp(a^k)}, \quad m = 0, \dots, N_y. \quad (16)$$

When NN_{CL} is exploited for prediction, the most likely class is selected as the one that best categorizes the processed measurements \mathbf{U} .

The model NN_{RG}^j addresses the regression task underlying the damage quantification problem, namely $\text{NN}_{\text{RG}}^j : \mathbf{U} \rightarrow \delta \in \mathbb{R}$, with $j = 1, \dots, N_y$. The estimated counterpart of δ is obtained as $\hat{\delta} = \text{NN}_{\text{RG}}^j(\mathbf{U})$. Hence, the regression models, one for each damageable region, map the vibration recordings \mathbf{U} associated with the j -th damage region, onto the estimated magnitude of the stiffness reduction taking place within the relative damage region. Since all NN_{RG}^j models are learned following the same procedure, the index j will be dropped in the following for ease of notation.

Since the space of digital states in the PGM is discrete, the outcomes of NN_{CL} and NN_{RG} are accommodated within the PGM by discretizing the range in which the damage level δ can take values in N_δ uniform intervals, thus resulting in $N_d = 1 + N_\delta N_y$ possible states. The same reasoning is followed to compute the confusion matrix encoding the ϕ_i^{NN} factor. In particular, ϕ_i^{NN} measures the offline performance of NN_{CL} and NN_{RG} in assimilating noisy FOM data to classify the digital state, among the N_d possible outcomes of D_i .

The models NN_{CL} and NN_{RG} are trained separately. The datasets dedicated to the training of NN_{CL} and NN_{RG} are derived from dataset D in Eq. (14) as follows. The dataset used to learn NN_{CL} is obtained from Eq. (14), as

$$D_{\text{CL}} = \{(\mathbf{U}_i, \mathbf{b}_i)\}_{i=1}^I. \quad (17)$$

The dataset used to learn NN_{RG} is derived from Eq. (14), as

$$D_{\text{RG}} = \{(\mathbf{U}_{i_{\text{RG}}}, \delta_{i_{\text{RG}}})\}_{i_{\text{RG}}=1}^{I_{\text{RG}}}, \quad (18)$$

where I_{RG} is the number of training instances in D_{RG} , all characterized by a structural damage within the same predefined region.

The set of weights and biases parametrizing NN_{CL} is denoted as Θ_{CL} . This is optimized minimizing the probabilistic categorical cross-entropy [37,52] \mathcal{L}_{CL} between the predicted and target class labels over D_{CL} :

$$\mathcal{L}_{\text{CL}}(\Theta_{\text{CL}}, D_{\text{CL}}) = -\frac{1}{I} \sum_{i=1}^I \sum_{m=0}^{N_y} b_i^m \log(\hat{b}_i^m), \quad (19)$$

which provides a measure of the distance between the discrete probability distribution describing \mathbf{b} , and its estimated counterpart $\hat{\mathbf{b}} = \text{NN}_{\text{CL}}(\mathbf{U})$.

The set of weights and biases Θ_{RG} parametrizing NN_{RG} is learned through the minimization of the following mean squared error loss function:

$$\mathcal{L}_{\text{RG}}(\Theta_{\text{RG}}, D_{\text{RG}}) = \frac{1}{I_{\text{RG}}} \sum_{i_{\text{RG}}=1}^{I_{\text{RG}}} (\delta_{i_{\text{RG}}} - \hat{\delta}_{i_{\text{RG}}})^2, \quad (20)$$




Fig. 4. L-shaped cantilever beam: details of synthetic recordings related to displacements $u_1(t), \dots, u_8(t)$, loading condition, and predefined damage regions $\Omega_1, \dots, \Omega_7$.

which provides a measure of the distance between the target magnitude of the stiffness reduction δ , and its approximated counterpart $\hat{\delta} = \text{NN}_{\text{RG}}(\mathbf{U})$.

The algorithmic description of the procedures and computations characterizing the preliminary offline phase of the proposed digital twinning framework is reported in Algorithm 2. The implementation details of the deep learning models are reported in the Appendix.

Algorithm 2 Preliminary offline phase – algorithmic description

Input: parametrization of the operational and damage conditions

PGM implementing the prediction graph

- 1: set up the physics-based numerical model of the structure to be monitored.
 - 2: assemble the snapshot matrix of the structural response via FOM analyses.
 - 3: compute the POD basis functions via singular value decomposition of the snapshots matrix.
 - 4: use the ROM to populate the training dataset D with vibration recordings at sensor location.
 - 5: use the recordings and labels in D to derive D_{CL} and D_{RG} .
 - 6: train the classification model NN_{CL} on D_{CL} and the regression models NN_{RG} on D_{RG} .
 - 7: test the generalization capabilities of NN_{CL} and NN_{RG} on noisy FOM data.
 - 8: compute the confusion matrix encoding the ϕ_i^{NN} factor.
 - 9: compute the control policy $\pi(D_i)$ by solving the planning problem induced by the PGM.
 - 10: **return** trained DL models, ϕ_i^{NN} factor, control policy $\pi(D_i)$.
-

3. Numerical experiments

This section demonstrates the proposed methodology for two test cases: an L-shaped cantilever beam and a railway bridge.

The FOM and ROM in Problem (11) and Problem (12) are implemented in the Matlab environment, using the redBKIT library [53]. The PGM framework for predictive digital twins is implemented in Python, using the pgmpy library [54]. All computations have been carried out on a PC featuring an AMD Ryzen™ 9 5950X CPU @ 3.4 GHz and 128 GB RAM. The NN architectures are implemented through the Tensorflow-based Keras API [55], and trained on a single Nvidia GeForce RTX™ 3080 GPU card.

3.1. L-shaped cantilever beam

The first test case deals with the L-shaped cantilever beam depicted in Fig. 4. The structure is made of two arms, each one having a length of 4 m, a width of 0.3 m, and a height of 0.4 m. The assumed mechanical properties are those of concrete: Young's modulus $E = 30$ GPa, Poisson's ratio $\nu = 0.2$, density $\rho = 2500$ kg/m³. The structure is excited by a distributed vertical load $q(t)$, acting on an area of (0.3×0.3) m² close to its tip. The load varies in time according to $q(t) = Q \sin(2\pi ft)$, with $Q \in [40, 80]$ kPa and $f \in [10, 60]$ Hz, respectively being the load amplitude and frequency. Following the setup described in Section 2, these parameters have a uniform distribution within their respective ranges.

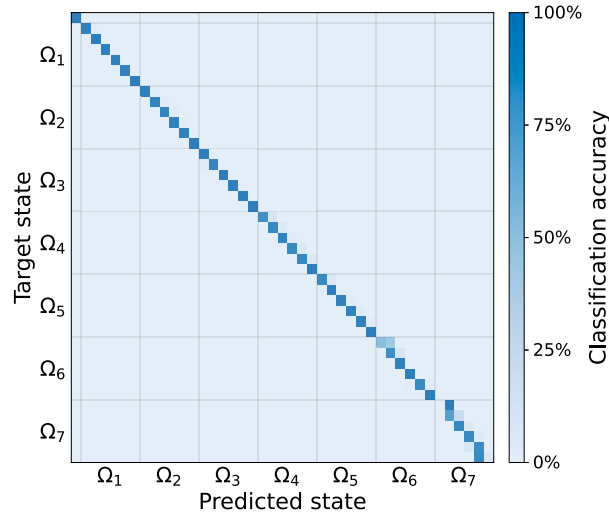


Fig. 5. L-shaped cantilever beam - Confusion matrix measuring the offline performance of the DL models in correctly categorizing the digital state. Results are reported in terms of classification accuracy, measuring how observational data are classified with respect to the ground truth digital state. Digital states are ordered first for damage location and then for damage level.

3.1.1. Dataset assembly

Synthetic displacement time histories \mathbf{U} are obtained in relation to $N_u = 8$ dofs along the bottom surface of the structure, to mimic the monitoring system depicted in Fig. 4. Each recording is provided for a time interval $(0, T = 1 \text{ s})$ with an acquisition frequency $f_s = 200 \text{ Hz}$. Recordings are corrupted with an additive Gaussian noise yielding a signal-to-noise ratio of 100.

In addition to the damage-free baseline condition, damage is simulated by considering $N_y = 7$ possible damage classes, each referring to a reduction of the material stiffness within a subdomain Ω_j , with $j = 1, \dots, N_y$, as depicted in Fig. 4. The stiffness reduction can occur with a magnitude $\delta \in [30\%, 80\%]$, and is held constant within the considered time interval.

The FOM is obtained with a finite element discretization using linear tetrahedral elements and resulting in $N_{FE} = 4659$ dofs. The basis matrix \mathbf{W} is obtained from a snapshot matrix \mathbf{S} , assembled through 400 evaluations of the FOM, at varying values of the input parameters $\boldsymbol{\mu} = (Q, f, y, \delta)^T$ sampled via Latin hypercube rule. By prescribing a tolerance $\epsilon = 10^{-3}$ on the fraction of energy content to be disregarded in the approximation, the order of the ROM approximation turns out to be $N_{RB} = 56$.

The dataset \mathcal{D} is built with $I = 10,000$ instances collected using the ROM. This is then employed to train NN_{CL} and NN_{RG} , as described in the previous section. In the absence of experimental data, the testing phase of NN_{CL} and of NN_{RG} is carried out through noise-corrupted FOM solutions. In particular, the asset is monitored by processing batches of $N_{\text{obs}} = 10$ noisy observations, relative to the same damage location y and damage magnitude δ , yet featuring varying operational conditions set by Q and f . As the health of the asset evolves over time, the DT assimilates a batch of noisy observations $\{\mathbf{U}_k\}_{k=1}^{N_{\text{obs}}}$ at each time step, to dynamically estimate the variation in the structural health parameters underlying the digital state.

3.1.2. Digital twin framework

The two structural health parameters within the digital state are $\mathbf{d} = (y, \delta)^T$. In order to accommodate the outcome of the DL models within the PGM and to compute the CPT encoding the ϕ_t^{NN} factor, the range in which the damage level δ can take values is discretized in $N_\delta = 6$ intervals $\{[30\%, 35\%], [35\%, 45\%], [45\%, 55\%], [55\%, 65\%], [65\%, 75\%], [75\%, 80\%]\}$, thus resulting in $N_d = 43$ possible digital states. The number of δ intervals and the width of each interval are chosen arbitrarily, and there are no restrictions in this respect. The resulting digital states are then sorted to follow the lexicographic order.

The confusion matrix reported in Fig. 5 measures the offline performance of NN_{CL} and NN_{RG} in assimilating noisy FOM data to classify the digital state, among the N_d possible outcomes of D_t . The (unknown) ground truth digital state is detected by the DL models with an overall classification accuracy of 93.61%. Moreover, it can be argued from the confusion matrix that most of the misclassifications are due to the damage scenarios related to a stiffness reduction within Ω_6 or within Ω_7 . This is a quite expected outcome since measurements closer to the clamped side are only marginally affected by the presence of damage close to the free end of the beam, thus yielding a smaller sensitivity of sensor recordings to damage. This confusion matrix then serves as the CPT encoding the ϕ_t^{NN} factor.

For the present case, we consider four possible control inputs, each provided with a CPT modeling the transition probability $p(D_{t+1}|D_t, U_t^A = u_t^A)$ from D_t to D_{t+1} after taking the action u_t^A , and collectively encoding the ϕ_t^{history} factor. These internal models of how structural health is expected to evolve do not reflect the prescribed ground truth evolution, which is unknown to the DT. The considered control inputs are the following:

- Do nothing (DN) action. There is no maintenance action planned in this case and the physical state will evolve according to a stochastic deterioration process.

- Minor imperfect maintenance (MI) action. A maintenance action is performed and the asset may be restored from its current condition to a healthier damage state. This can be traced back to, e.g., patching and sealing cracked surfaces, rectifying and replacing expansion joints, or tightening/replacing loose/missing knot bolts for steel members.
- Major imperfect maintenance (MA) action. A maintenance action is performed and the asset may be restored from its current condition to a healthier damage state, with a higher probability of improvements than in the previous case. This can be traced back to, e.g., repairing heavily damaged slabs, piers, and steel members, and retrofitting compromised structural elements.
- Perfect maintenance (PM) action. A maintenance action is performed and the asset is restored from its current condition to the damage-free state. This can be traced back to the replacement of excessively compromised structural elements.

3.1.3. Results: two available actions

We first illustrate the DT capabilities to assimilate observational data and track the structural health evolution, by restricting the available actions to DN and PM. The (unknown) ground truth evolution of structural health varies depending on the most recently applied control input, which can be either DN or PM. In the absence of maintenance, the physical state evolves following a deterioration process. We prescribe a (simulated) stochastic degradation process that monotonically deteriorates the structural health. The degradation process features a probability of damage inception ($y \neq 0$) equal to 0.5. Damage may develop in any of the predefined regions with $\delta = 30\%$, and then propagate with δ increments sampled from a Gaussian probability density function (pdf) centered at 1.5% and featuring a standard deviation equal to 1% (negative increments are rounded to zero). The effect of a PM action is simulated by restoring the physical state to its undamaged configuration. At each time step during the operation, new observational data are simulated according to the (unknown) ground truth structural health and the most recently enacted control input. The DT assimilates the data and estimates the digital state, eventually suggesting the next control input to enact. Note that the prescribed trajectory of the structural health parameters is arbitrarily chosen to fully display the capabilities of the DT. Nevertheless, the DT would be equally capable of tracking the structural health evolution also considering either more or less aggressive degradation processes.

The state transition model encoding ϕ_t^{history} is conditioned on the most recently issued control input. The transition probability $p(D_{t+1}|D_t, U_t^A = u_t^A)$ from D_t to D_{t+1} associated with the DN action assumes that damage may start in any subdomain Ω_j , with $j = 1, \dots, N_y$, with probability 0.05, and then grow to the next δ interval with the same probability. The transition model assumed for the PM action instead maps the D_t belief to a belief D_{t+1} associated with a damage-free condition, independently of the current condition. The corresponding CPTs are transition matrices, where the diagonal entries represent the probability of staying in the same state. The lower-left and upper-right triangles are associated with the probabilities of the system of deteriorating and improving its condition, respectively. Therefore, the DN transition matrix is a lower-left triangular matrix, with the highest probability assigned to remaining in the same state, consistent with what is expected for the deterioration of civil structures. The transition to the next δ interval is the second most likely transition, while improvements have a zero probability. Once the structure has reached the last δ interval, it remains in this condition with a probability equal to 1. In contrast, the PM transition matrix is an upper-right triangular matrix with probabilities equal to 1 in the first row.

At each time step, the DT selects a control input $u_t^A \in \{\text{DN}, \text{PM}\}$ to be enacted on the asset. Taking a DN action yields a positive reward, but also gives the chance of worsening the asset's structural health. On the other hand, the PM action responds to the degrading structural health, yet yields a negative reward. The computation of the costs associated with the health state and control inputs encapsulates the evaluation of the ϕ_t^{reward} factor quantifying the performance of the asset. In particular, the two reward functions in Eq. (9) are defined as:

$$R_t^{\text{control}}(u_t^A) = \begin{cases} +12, & \text{if } u_t^A = \text{DN}, \\ -20, & \text{if } u_t^A = \text{PM}, \end{cases} \quad R_t^{\text{health}}(d_t) = \begin{cases} +0.1, & \text{if } y = 0, \\ -\exp(6\delta/5), & \text{if } y \in \{1, 2, 3, 4\}, \\ -\exp(\delta), & \text{if } y \in \{5, 6, 7\}, \end{cases} \quad (21)$$

where R_t^{control} targets the cost assigned to each control input and R_t^{health} measures the cost associated with the structural health state. These non-dimensional rewards represent indicative values the decision-maker is charged due to the condition of the structure. Although these values are not based on real data, actual values are not usually hard to find. State agencies and companies provide lists with services and costs [56]. The three cases in R_t^{health} distinguish between the absence of damage, the presence of damage within the harm closed to the clamped side, and the presence of damage far from the clamp, respectively. Note how these penalize the progressive deterioration of the structural health as a function of δ . R_t^{health} can resemble a variety of aspects, like reduction in the level of service due to deterioration, working accidents, structural reliability, and structural failure probability [56].

During the offline phase, we solve the planning problem induced by the PGM to compute the control policy $\pi(D_t)$, which maps the digital state belief to actions and encodes the control factor ϕ_t^{control} . The optimization of $R_t(U_t, D_t)$ is carried out as described in Section 2.1, assuming a discount factor $\gamma = 0.95$ and a weighting factor $\alpha = 2$. The computed control policy recommends that the asset operates until when $\delta \in [65\%, 75\%]$ and $\delta \in [75\%, 80\%]$, respectively if $y \in \{1, 2, 3, 4\}$ and if $y \in \{5, 6, 7\}$, at which point it should be repaired.

Fig. 6 depicts a simulated online phase of the DT up to time step $t_c = 50$. Results are reported in terms of the (unknown) ground truth digital state, and the corresponding DT estimate after assimilating the observational data. The graphs report the evolution of the digital state only for the damaged regions, nevertheless, damage can potentially affect all $\Omega_1, \dots, \Omega_7$ predefined damageable regions. The DT proves capable of accurately tracking the digital state evolution with relatively low uncertainty. The corresponding estimation of the control inputs is reported in the bottom part of the figure, demonstrating that the DT is able to promptly suggest the PM action within one time step of when the (unknown) ground truth structural health demands it.

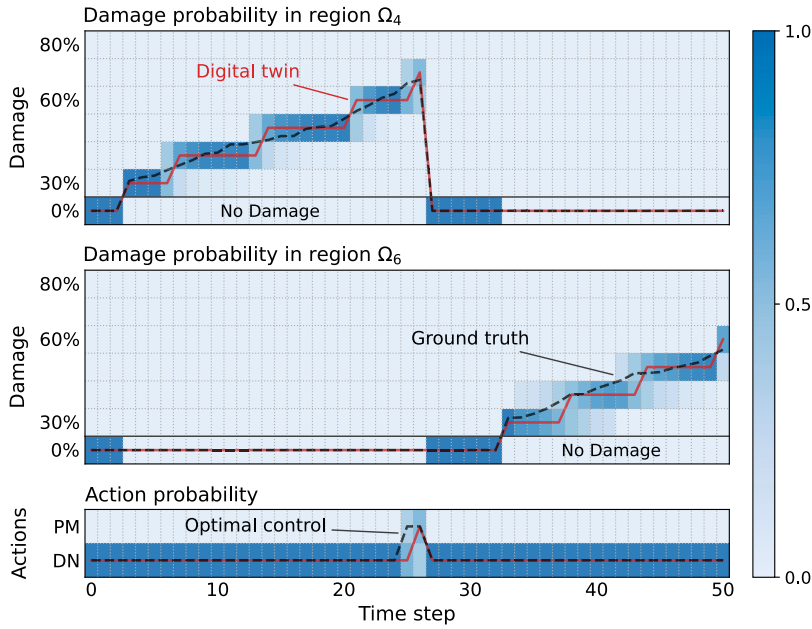


Fig. 6. L-shaped cantilever beam - Online phase of the digital twin framework with two possible actions: DN (do nothing), and PM (perfect maintenance). Probabilistic and best point estimates of: (top) digital state evolution against the ground truth digital state; (bottom) control inputs informed by the digital twin, against the optimal control input under ground truth. In the top panels the background color corresponds to $p(D_t|D_{t-1}, D_t^{DN}, U_{t-1}^A = u_{t-1}^A)$. In the bottom panel it corresponds to $p(U_t|D_t)$.




Fig. 7. L-shaped cantilever beam - Digital twin future predictions with two possible actions: DN (do nothing), and PM (perfect maintenance). The starting time is $t_c = 50$. In the top panel the probability $p(D_t|D_{t-1}, U_{t-1})$ relates to the amount of damage in Ω_6 . In the bottom panel it corresponds to $p(U_t|D_t)$.

Fig. 7 depicts the predicted evolution of the digital state and of the corresponding informed control inputs, starting from $t_c = 50$. The prediction horizon is extended over 20 time steps in the future so that $t_p = t_c + 20$. The DT prediction engine informs about the expected future degradation of the structural health, allowing to plan future interventions.

3.1.4. Results: four available actions

We now consider all four possible control inputs. We prescribe a stochastic degradation process with a probability of damage inception ($\gamma \neq 0$) equal to 0.5. Damage may develop in any of the predefined regions with damage level sampled from a uniform distribution $\delta \in [30\%, 70\%]$, and then propagate as in the previous case. This more aggressive degradation process is used to spot in a few time steps the effectiveness of the decision-making capabilities of the DT. The effect of the MI and MA actions on the asset is simulated according with stochastic repair processes, for which the structural health is forced to improve. The effect of a MI action is simulated with δ decrements sampled from a Gaussian pdf centered at -12.5% and featuring a standard deviation equal to 1%,

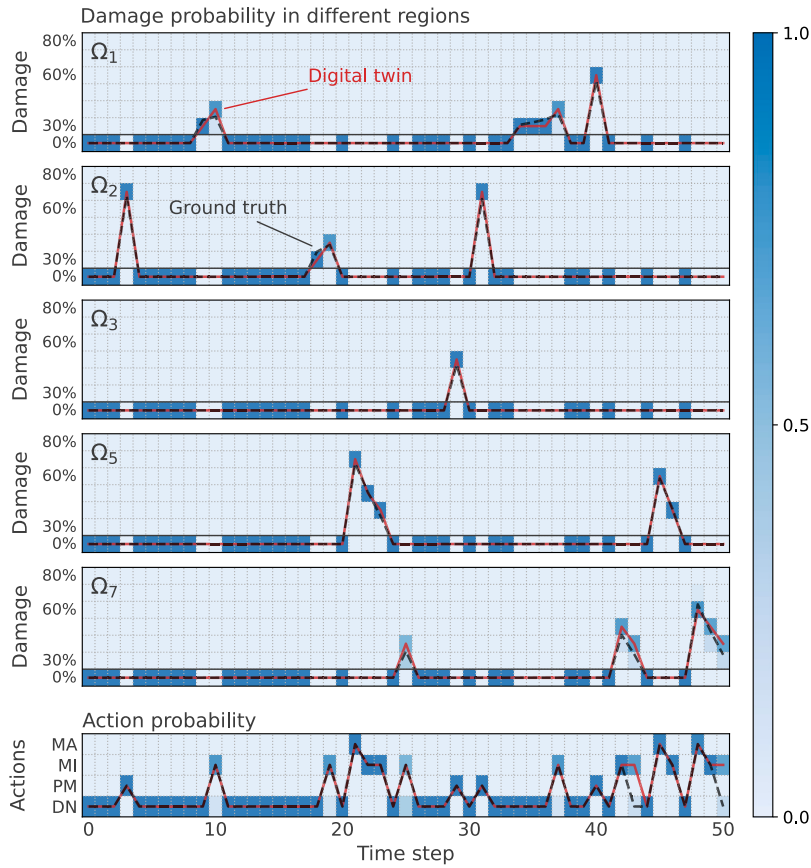


Fig. 8. L-shaped cantilever beam - Online phase of the digital twin framework with four possible actions: DN (do nothing), PM (perfect maintenance), MI (minor imperfect maintenance), and MA (major imperfect maintenance). Probabilistic and best point estimates of: (top) digital state evolution against the ground truth digital state; (bottom) control inputs informed by the digital twin, against the optimal control input under ground truth. In the top panels the background color corresponds to $p(D_t|D_{t-1}, D_t^{NN}, U_{t-1}^A = u_{t-1}^A)$. In the bottom panel it corresponds to $p(U_t|D_t)$.

while the effect of a MA action is modeled with δ decrements sampled from a Gaussian pdf centered at -17.5% and featuring a standard deviation equal to 1%. In both cases, the damage-free condition is assumed to be recovered if the resulting structural state features $\delta < 30\%$.

The transition model $p(D_{t+1}|D_t, U_t^A = u_t^A)$ associated with the MI action assumes no improvement in the structural health with probability 0.1, improvement of one δ interval with probability 0.75, and improvement of two δ intervals with probability 0.15. The resulting CPT is an upper-right triangular transition matrix, as deterioration from any state upon a repair action is assumed to have zero probability. The highest probability is assigned to improvements of one δ interval, followed by improvements of two δ intervals. There is also a lower probability of remaining in the same deteriorated state, which reflects a failed maintenance. Similarly, the MA action assumes no improvement with probability 0.05, improvement of one δ interval with probability 0.3, improvement of two δ intervals with probability 0.4, and improvement of three δ intervals with probability 0.25. In this case, the highest probability is assigned to improvements of two δ intervals, followed by improvements of one δ intervals, three δ intervals, and finally, the lowest probability is associated with the possibility of a failed maintenance.

The two reward functions in Eq. (9) are chosen as:

$$R_t^{\text{control}}(u_t^A) = \begin{cases} +12, & \text{if } u_t^A = \text{DN}, \\ -20, & \text{if } u_t^A = \text{PM}, \\ -8, & \text{if } u_t^A = \text{MI}, \\ -15, & \text{if } u_t^A = \text{MA}, \end{cases} \quad R_t^{\text{health}}(d_t) = \begin{cases} +0.1, & \text{if } y = 0, \\ -\exp(5\delta), & \text{if } y \in \{1, 2, 3, 4\}, \\ -\exp(4\delta), & \text{if } y \in \{5, 6, 7\}. \end{cases} \quad (22)$$

We assume a discount factor $\gamma = 0.95$, and a weighting factor $\alpha = 2.5$. The resulting control policy $\pi(D_t)$ recommends that the asset should operate until when $\delta \in [30\%, 35\%]$, after which: if $y \in \{1, 2, 3, 4\}$, a MI action should be performed when $\delta \in [35\%, 45\%]$, and a PM action should be performed when $\delta > 45\%$; while, if $y \in \{5, 6, 7\}$, the MI and MA actions should be performed, respectively when $\delta \in [35\%, 55\%]$ and when $\delta \in [55\%, 75\%]$, and a PM action should be performed when $\delta > 75\%$.

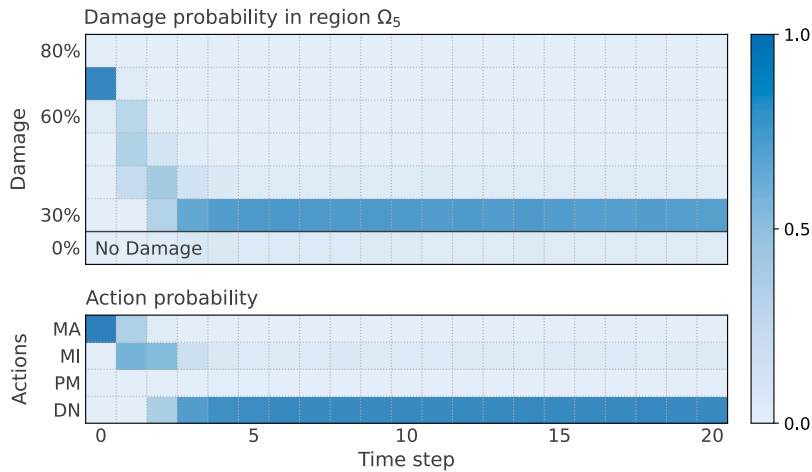


Fig. 9. L-shaped cantilever beam - Digital twin future predictions with four possible actions: DN (do nothing), PM (perfect maintenance), MI (minor imperfect maintenance), and MA (major imperfect maintenance). The starting time is $t_c = 21$. In the top panel the probability $p(D_t|D_{t-1}, U_{t-1})$ relates to the amount of damage in Ω_5 . In the bottom panel it corresponds to $p(U_t|D_t)$.

Fig. 8 depicts a simulated online phase of the DT up to $t_c = 50$. The DT accurately tracks the digital state evolution and timely suggests the appropriate control inputs most of the time. In particular, the DT proposes the optimal control input, except for the time steps $t = 43$ and $t = 50$ featuring a sub-optimal action. In both cases, a MI action is proposed in place of a DN, because the DT estimates a $\delta \in [35\%, 45\%]$ instead of a $\delta \in [30\%, 35\%]$ related to a stiffness reduction within Ω_7 . This is in line with what was observed in the confusion matrix of **Fig. 5**, due to the limited sensitivity of recordings to damage scenarios affecting the terminal region of the beam. This peculiar type of misclassification turns out to be the most pathological in the confusion matrix and is therefore capable of potentially spoiling the assimilation of observational data. Nevertheless, the DT reverts to correctly tracking the structural health of the asset within one time step.

Fig. 9 depicts the predicted evolution of the digital state and control inputs, from $t_c = 21$ and over 20 time steps in the future. The DT prediction correctly suggests taking with high probability a MA action, followed by two MI actions, and accordingly predicts the corresponding evolution of the structural health. Comparing the DT prediction with what is effectively experienced during the online phase (see **Fig. 8**), note how the DT prediction closely resembles the actual evolution of the digital state and control inputs. This is a remarkable result in terms of DT prediction capabilities, since the DT is not aware of the future values of the structural health parameters, and the relative transition models do not match their real (stochastic) evolution.

3.2. Railway bridge

The second case study concerns the railway bridge depicted in **Fig. 10**. It is an integral concrete portal frame bridge located along the Bothnia line in the Swedish suburbs of Hörnefors. It features a span of 15.7 m, a free height of 4.7 m, and a width of 5.9 m (edge beams excluded). The thickness of the structural elements is 0.5 m for the deck, 0.7 m for the frame walls, and 0.8 m for the wing walls. The bridge is founded on two plates connected by stay beams and supported by pile groups. The concrete is of class C35/45, whose mechanical properties are: $E = 34$ GPa, $\nu = 0.2$, $\rho = 2500$ kg/m³. The superstructure consists of a single track with sleepers spaced 0.65 m apart, resting on a ballast layer 0.6 m deep, 4.3 m wide and featuring a density $\rho_B = 1800$ kg/m³. The geometrical and mechanical modeling data have been adapted from former research activities on the relevant soil–structure interaction, see [57,58].

The bridge is subjected to the transit of Gröna Tåget trains type, at a speed $v \in [160, 215]$ km/h. Only trains composed of two wagons are considered, thus characterized by 8 axles, each one carrying a mass $\psi \in [16, 22]$ ton. The corresponding load model is described in [38], and consists of 25 equivalent distributed forces transmitted by the sleepers to the deck through the ballast layer with a slope 4:1, according with Eurocode 1 [59].


3.2.1. Dataset assembly

Synthetic displacement time histories U are obtained from $N_u = 10$ sensors deployed as depicted in **Fig. 11**. Each recording is provided for a time interval $(0, T = 1.5$ s) with an acquisition frequency $f_s = 400$ Hz. This setting allows to record train passages at the lowest speed of 160 km/h, and properly catches the structural response at the maximum speed of 215 km/h. Recordings are corrupted with an additive Gaussian noise yielding a signal-to-noise ratio of 120.

In addition to the undamaged condition, the presence of damage in the structure is accounted for using a localized stiffness reduction that can take place within $N_y = 6$ predefined subdomains Ω_j , with $j = 1, \dots, N_y$, as depicted in **Fig. 11**. The stiffness reduction can occur with a magnitude $\delta \in [30\%, 80\%]$, and is kept fixed while a train travels across the bridge.



Fig. 10. Hörnefors railway bridge.

Fig. 11. Railway bridge: details of synthetic recordings related to displacements $u_1(t), \dots, u_{10}(t)$, and predefined damage regions $\Omega_1, \dots, \Omega_6$.

The FOM features $N_{FE} = 17,292$ dofs, resulting from a finite element discretization with an element size of 0.80 m and a reduced size of 0.15 m for the deck, to enable a smooth propagation of the traveling load. The presence of the ballast layer is accounted for through an increased density for the deck and for the edge beams. The embankments are accounted for through distributed springs, modeled as a Robin mixed boundary condition (with elastic coefficient $k_{robin} = 10^8 \text{ N/m}^3$) applied on the surfaces facing the ground. The structural dissipation is modeled by means of a Rayleigh's damping matrix, assembled to account for a 5% damping ratio on the first two structural modes.

The ROM is obtained from a snapshot matrix S , assembled through 400 evaluations of the FOM for different values of parameters $\mu = (v, \psi, y, \delta)^T$. By setting the error tolerance to $\epsilon = 10^{-3}$, $N_{RB} = 133$ POD modes are to be considered.

The training dataset D is built with $I = 10,000$ instances collected using the ROM. Also in this case, the testing phase of NN_{CL} and of NN_{RG} is carried out considering noisy FOM solutions. The monitoring of the asset is then simulated by assimilating $N_{obs} = 1$ noisy observations at each time step. As the structural health of the bridge evolves over time, the DT estimates the variation in the structural health parameters every time a train travels across the bridge.

3.2.2. Digital twin framework

As in the previous case, the two structural health parameters within the digital state are $d = (y, \delta)^T$. The range in which the damage level δ can take values is discretized in $N_\delta = 6$ intervals. The resulting $N_d = 37$ possible digital states are sorted first for damage location and then for damage level.

The confusion matrix measuring the offline performance of NN_{CL} and of NN_{RG} in correctly categorizing the digital state is reported in Fig. 12. The ground truth digital state is detected with an overall classification accuracy of 91.39%. In this case, the majority of misclassifications are due to confusing adjacent digital states relative to the same damage location, thus yielding a tridiagonal band matrix.

For the present case, we consider the following three possible control inputs:

- Do nothing (DN) action. There is no maintenance action planned in this case and the physical state will evolve according to a stochastic deterioration process.

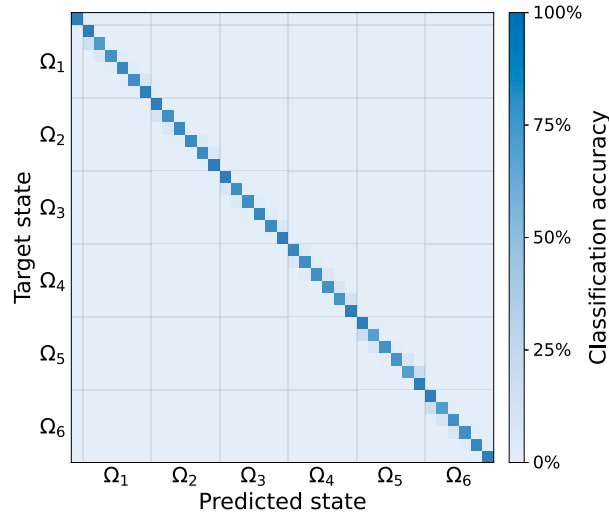


Fig. 12. Railway bridge - Confusion matrix measuring the offline performance of the DL models in correctly categorizing the digital state. Results are reported in terms of classification accuracy, measuring how observational data are classified with respect to the ground truth digital state. Digital states are ordered first for damage location and then for damage level.

- Perfect maintenance (PM) action. A maintenance action is performed and the asset is restored from its current condition to the damage-free state.
- Restrict operational conditions (RE) action. The operational conditions of the bridge are restricted by allowing only lightweight trains, carrying less than 18 ton per axle, to travel across the bridge. Such a restriction results in a lower deterioration rate, but also yields a lower revenue generated by the infrastructure.

In the cases where the most recently issued control input is either DN or RE, the physical state undergoes a degradation process that monotonically deteriorates the structural health. When operational conditions are not restricted, we prescribe a stochastic degradation process featuring a probability of damage inception ($y \neq 0$) equal to 0.5. Damage may develop in any of the predefined regions with damage level sampled from a uniform distribution $\delta \in [30\%, 35\%]$, and then propagate with δ increments sampled from a Gaussian pdf centered at 1.5% and featuring a standard deviation equal to 1% (negative increments are rounded to zero). When the operations are restricted and only lightweight trains are allowed to travel across the bridge, we instead assume a probability of damage inception equal to 0.25. In this eventuality, damage may develop with damage level sampled from a uniform distribution $\delta \in [30\%, 35\%]$, and then propagate with δ increments sampled from a Gaussian pdf centered at 0.95% and featuring a standard deviation equal to 0.5%. The resulting trajectory of the structural health parameters is intended to represent periods of gradual degradation in the structural health, as well as sudden changes due to discrete damage events. Also in this case, the effect of a PM action is simulated by restoring the physical state to its undamaged configuration.

The transition model $p(D_{t+1} | D_t, U_t^A = u_t^A)$ associated with the DN action assumes that damage may start in any subdomain Ω_j , with $j = 1, \dots, N_j$, with probability 0.1, and then grow to the next δ interval with the same probability. For the transition model associated with the RE action, this probability is assumed to decrease to 0.03. The CPTs associated with the DN and RE actions are therefore lower-left triangular transition matrices. The highest probability is assigned to remaining in the same state, followed by the transition to the next δ interval, with zero probability of improvements. The transition model assumed for the PM action instead maps the D_t belief to a belief D_{t+1} associated with a damage-free condition, independently of the current condition. The CPT associated with the PM action is therefore an upper-right triangular transition matrix with probabilities equal to 1 in the first row.

In this case, the two reward functions in Eq. (9) are chosen as:

$$R_t^{\text{control}}(u_t^A) = \begin{cases} +30, & \text{if } u_t^A = \text{DN}, \\ -250, & \text{if } u_t^A = \text{PM}, \\ +27, & \text{if } u_t^A = \text{RE}, \end{cases} \quad R_t^{\text{health}}(d_t) = \begin{cases} +0, & \text{if } y = 0, \\ -\exp(5\delta) + 4, & \text{if } y \neq 0, \\ -250, & \text{if } \delta \geq 79\%, \end{cases} \quad (23)$$

where the last contribution in R_t^{health} penalizes excessively compromised structural states with a significantly negative reward.

3.2.3. Results

During the offline phase, we solve the planning problem in Eq. (8) by assuming a discount factor $\gamma = 0.90$, and a weighting factor $\alpha = 1$. The resulting control policy $\pi(D_t)$ recommends that the asset operates in ordinary conditions until when $\delta \in [30\%, 35\%]$, after which point it should fall back to the more conservative RE regime in order to minimize further degradation. Once reached $\delta \geq 65\%$, the bridge should be finally repaired.




Fig. 13. Railway bridge - Online phase of the digital twin framework with three possible actions: DN (do nothing), PM (perfect maintenance), and RE (restrict operational conditions). Probabilistic and best point estimates of: (top) digital state evolution against the ground truth digital state; (bottom) control inputs informed by the digital twin, against the optimal control input under ground truth. In the top panels the background color corresponds to $p(D_t|D_{t-1}, D_t^{NN}, U_{t-1}^A = u_{t-1}^A)$. In the bottom panel it corresponds to $p(U_t|D_t)$.




Fig. 14. Railway bridge - Digital twin future predictions with three possible actions: DN (do nothing), PM (perfect maintenance), and RE (restrict operational conditions). The starting time is $t_c = 5$. In the top panel the probability $p(D_t|D_{t-1}, U_{t-1})$ relates to the amount of damage in Ω_5 . In the bottom panel it corresponds to $p(U_t|D_t)$.

Fig. 13 reports a sample simulation of the DT online phase up to time step $t_c = 60$. The DT correctly tracks the digital state with relatively low uncertainty. Damage initially develops within Ω_5 , and the DT follows its evolution with a limited delay of at most two time steps, with respect to the ground truth, due to the need of updating the relative prior belief from the previous time steps. The RE action is suggested as soon as the DT estimates a $\delta \in [35\%, 65\%]$, after which point the DT keeps on tracking the structural health parameters evolving with a lower deterioration rate. A PM action is finally suggested due to an excessively compromised structural state. A similar behavior can be observed for the following damage scenario affecting Ω_6 .

Fig. 14 reports the predicted evolution of the digital state and control inputs, from $t_c = 5$ and over 20 time steps in the future. The DT predicts the expected degradation of the structural health according to the transition model associated with the DN action, before predicting to take a RE action with relatively high probability after a few time steps. The DT prediction is close to what is effectively experienced online (see Fig. 13). However, besides having the estimated digital state two time steps behind the ground

truth value, the prediction is also too optimistic in terms of deterioration rate, which suggests the use of a more refined transition model.

4. Conclusions

In this work we have proposed a predictive digital twin approach to the health monitoring, maintenance, and management planning of civil structures, to advance condition-based and predictive maintenance practices. The presented strategy relies upon a probabilistic graphical model inspired by [18]. This framework is used to encode the asset-twin coupled dynamical system, the relevant end-to-end information flow via observational data (physical to digital) and control inputs (digital to physical), and its evolution over time, all with quantified uncertainty. The assimilation of observational data is carried out with deep learning models, leveraging the capabilities of convolutional layers to automatically select and extract damage-sensitive features from raw vibration recordings. The structural health parameters comprising the digital state are used to capture the variability of the physical asset. They are continually updated in a sequential Bayesian inference fashion, according to control-dependent transition dynamics models describing how the structural health is expected to evolve. The updated digital state is eventually exploited to predict the future evolution of the physical system and the associated uncertainty. This enables predictive decision-making about maintenance and management actions.

The computational procedure takes advantage of a preliminary offline phase which involves: (i) using physics-based numerical models and reduced order modeling, to overcome the lack of experimental data for civil applications under varying damage and operational conditions while populating the datasets for training the deep learning models; (ii) learning the health-dependent control policy to be applied at each time step of the online phase, to map the belief over the digital state onto actions feeding back to the physical asset.

The proposed strategy has been assessed against the simulated monitoring of an L-shaped cantilever beam and a railway bridge. In the absence of experimental data, the tests have been carried out considering high-fidelity simulation data, corrupted with an additive Gaussian noise. The obtained results have proved the digital twin capabilities of accurately tracking the digital state evolution under varying operational conditions, with relatively low uncertainty. The framework is also able to promptly suggest the appropriate control input, within at most two time steps of when the (unknown) ground truth structural health demands it.

Although the capabilities of health-aware digital twins are showcased in the specific context of monitoring the structural integrity of civil structures to advance predictive maintenance practices, the applicability of the presented framework is general. Indeed, the proposed framework can be adapted for various types of structures and engineering systems by adjusting the components within the dynamic Bayesian network to align with the specific characteristics of the problem at hand. The solution to the inverse problem (if any) can be estimated by assimilating available observational data using methods other than deep neural networks, for instance through Markov chain Monte Carlo sampling algorithms. Similarly, the state transition models are closely tied to the employed parametrization of the digital state and the availability of historical data. The same applies to the available control inputs, which are likely to vary for different structures, such as those in mechanical or aerospace systems, and the method chosen for solving the associated planning problem. Additionally, the graph topology can be easily reorganized to adapt to situations where observational data are not acquired after issuing a control input, or when control inputs are issued with a different frequency than that governing the digital twin update.

Future research lines will investigate the ability of the digital twin to update the transition dynamics models by learning from previous data. As suggested by the railway bridge case study, this will allow for a more accurate prediction of the expected evolution of the digital state, thus enabling predictive decision-making better tailored to the monitored asset. Another aspect of interest concerns solving the planning problem induced by the probabilistic model using reinforcement learning algorithms, capable of taking into account a finite planning horizon representing the design lifetime of the asset.

Code availability

The implementation code used for the experiments presented in Section 3 is available in the public repository `digital-twin-SHM` [28]. The code implements the proposed digital twin framework and can be used to generate the graphs of digital state estimation and prediction reported in this paper. The DL models trained according to the implementation details reported in the [Appendix](#) are also made available in the same repository.

Declaration of competing interest

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

Data availability

The observational data used to run the experiments presented in Section 3 are available in the public repository `digital-twin-SHM` [28]. The Matlab library for finite element simulation and reduced-order modeling of partial differential equations employed to generate these data is available in the public repository `Redbkit` [53].

Acknowledgments

Matteo Torzoni acknowledges the financial support from the Politecnico di Milano, Italy through the interdisciplinary Ph.D. Grant “Physics-informed deep learning for structural health monitoring”. Marco Tezzele and Karen E. Willcox acknowledge support from the “NASA University Leadership Initiative” under Cooperative Agreement 80NSSC21M0071. Andrea Manzoni acknowledges the project “Dipartimento di Eccellenza 2023–2027”, funded by MUR, Italy, and the project “FAIR (Future Artificial Intelligence Research)”, funded by the NextGenerationEU program within the PNRR-PE-AI scheme (M4C2, Investment 1.3, Line on Artificial Intelligence).

Appendix. Implementation details

In this Appendix, we discuss the implementation details of the DL models described in Section 2.3. The architectures, as well as the relevant hyperparameters and training options, have been chosen through a preliminary study, aimed at minimizing \mathcal{L}_{CL} and \mathcal{L}_{RG} , while retaining the generalization capabilities of NN_{CL}^j and of NN_{RG}^j , with $j = 1, \dots, N_y$. Since all NN_{RG}^j models share the same architecture, the index j will be dropped in the following for ease of notation.

In the present work, NN_{CL} and NN_{RG} are set as 12-layers DL models, whose architecture is outlined in Table A.1a and in Table A.2a, respectively. NN_{CL} and NN_{RG} feature a damage-sensitive feature extractor required to be insensitive to transformations in the input not related to damage. This is implemented through the composition of three one-dimensional (1D) convolutional units. Convolutional layers naturally embed good relational inductive biases such as locality and translation equivariance [52,60], which prove highly effective to detect time correlations within and across time series. The resulting sparse connectivity and parameter sharing also make them computationally efficient. Each convolutional unit consists of a convolutional layer, followed by a Tanh activation function, max pooling, and dropout. The extracted features are expected to be sensitive to the presence of damage, but insensitive to measurement noise and operational variability. The extracted features are then reshaped through a flatten layer and run through a stack of three fully-connected layers: the first two are Tanh-activated, while the output layer of NN_{CL} is Softmax-activated, and the output layer of NN_{RG} has no activation function.

Table A.1

NN_{CL} - (a) employed architecture, and (b) selected hyperparameters and training options.

(a)				(b)	
Layer	Output shape	Activation	Input		
0 - Input	(B_{CL}, L, N_y)	None	–	Convolution kernel size:	25, 13, 7
1 - Conv1D	$(B_{CL}, L, 32)$	Tanh	0	Dropout rate:	5%
2 - MaxPooling1D	$(B_{CL}, L/2, 32)$	None	1	Weight initializer:	Xavier
3 - Dropout	$(B_{CL}, L/2, 32)$	None	2	L^2 regularization rate:	$\lambda_{CL} = 10^{-3}$
4 - Conv1D	$(B_{CL}, L/2, 64)$	Tanh	3	Optimizer:	Adam
5 - MaxPooling1D	$(B_{CL}, L/4, 64)$	None	4	Batch size:	$B_{CL} = 32$
6 - Dropout	$(B_{CL}, L/4, 64)$	None	5	Initial learning rate:	$\eta_{CL} = \{10^{-3}, 10^{-4}\}$
7 - Conv1D	$(B_{CL}, L/4, 32)$	Tanh	6	Allowed epochs:	250
8 - MaxPooling1D	$(B_{CL}, L/8, 32)$	None	7	Learning schedule:	$\frac{4}{5}$ cosine decay
9 - Dropout	$(B_{CL}, L/8, 32)$	None	8	Weight decay:	0.05
10 - Flatten	$(B_{CL}, 4L)$	None	9	Early stop patience:	15 epochs
11 - Dense	$(B_{CL}, 64)$	Tanh	10	Train-val split:	80 : 20
12 - Dense	$(B_{CL}, 16)$	Tanh	11		
13 - Dense	$(B_{CL}, N_y + 1)$	Softmax	12		

Table A.2

NN_{RG} - (a) employed architecture, and (b) selected hyperparameters and training options.

(a)				(b)	
Layer	Output shape	Activation	Input		
0 - Input	(B_{RG}, L, N_u)	None	–	Convolution kernel size:	25, 13, 7
1 - Conv1D	$(B_{RG}, L, 32)$	Tanh	0	Dropout rate:	10%
2 - MaxPooling1D	$(B_{RG}, L/2, 32)$	None	1	Weight initializer:	Xavier
3 - Dropout	$(B_{RG}, L/2, 32)$	None	2	L^2 regularization rate:	$\lambda_{RG} = 10^{-3}$
4 - Conv1D	$(B_{RG}, L/2, 64)$	Tanh	3	Optimizer:	Adam
5 - MaxPooling1D	$(B_{RG}, L/4, 64)$	None	4	Batch size:	$B_{RG} = 32$
6 - Dropout	$(B_{RG}, L/4, 64)$	None	5	Initial learning rate:	$\eta_{RG} = \{10^{-3}, 10^{-4}\}$
7 - Conv1D	$(B_{RG}, L/4, 32)$	Tanh	6	Allowed epochs:	250
8 - MaxPooling1D	$(B_{RG}, L/8, 32)$	None	7	Learning schedule:	$\frac{4}{5}$ cosine decay
9 - Dropout	$(B_{RG}, L/8, 32)$	None	8	Weight decay:	0.05
10 - Flatten	$(B_{RG}, 4L)$	None	9	Early stop patience:	15 epochs
11 - Dense	$(B_{RG}, 64)$	Tanh	10	Train-val split:	80 : 20
12 - Dense	$(B_{RG}, 16)$	Tanh	11		
13 - Dense	$(B_{RG}, 1)$	None	12		

Using the Xavier's weight initialization [61], NN_{CL} and NN_{RG} are trained by minimizing the following loss functions, respectively:

$$\mathcal{L}_{\text{CL}}^R(\Theta_{\text{CL}}, D_{\text{CL}}) = \mathcal{L}_{\text{CL}}(\Theta_{\text{CL}}, D_{\text{CL}}) + \lambda_{\text{CL}} \|\Theta_{\text{CL}}\|_2^2, \quad (\text{A.1})$$

$$\mathcal{L}_{\text{RG}}^R(\Theta_{\text{RG}}, D_{\text{RG}}) = \mathcal{L}_{\text{RG}}(\Theta_{\text{RG}}, D_{\text{RG}}) + \lambda_{\text{RG}} \|\Theta_{\text{RG}}\|_2^2, \quad (\text{A.2})$$

where λ_{CL} and λ_{RG} denote the L^2 regularization rate over the relative model parameters Θ_{CL} and Θ_{RG} . The loss functions $\mathcal{L}_{\text{CL}}^R$ and $\mathcal{L}_{\text{RG}}^R$ are minimized using the first-order stochastic gradient descent optimizer Adam [62], for a maximum of 250 allowed epochs. The corresponding learning rates η_{CL} and η_{RG} are initially set to $\{10^{-3}, 10^{-4}\}$, and decreased for 4/5 of the allowed training steps using a cosine decay schedule with weight decay equal to 0.05. The optimization is carried out considering an 80:20 splitting ratio of the dataset for training and validation purposes, with 20% of the data randomly taken and set aside to monitor the learning process. We use an early stopping strategy to interrupt learning, whenever the loss function value attained on the validation set does not decrease for a prescribed number of patience epochs in a row. The hyperparameters and training options for NN_{CL} and for NN_{RG} are reported in Table A.1b and in Table A.2b, respectively.

References

- [1] A. Thelen, X. Zhang, O. Fink, Y. Lu, S. Ghosh, B.D. Youn, M.D. Todd, S. Mahadevan, C. Hu, Z. Hu, A comprehensive review of digital twin – part 1: modeling and twinning enabling technologies, *Struct. Multidiscip. Optim.* 65 (2022) 354, <http://dx.doi.org/10.1007/s00158-022-03425-4>.
- [2] E.J. Tuegel, A.R. Ingraffea, T.G. Eason, S.M. Spottswood, Reengineering aircraft structural life prediction using a digital twin, *Int. J. Aerosp. Eng.* 2011 (2011) 154798, <http://dx.doi.org/10.1155/2011/154798>.
- [3] E. Tuegel, The airframe digital twin: Some challenges to realization, in: 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference, 2012, <http://dx.doi.org/10.2514/6.2012-1812>.
- [4] M. Grieves, J. Vickers, Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems, in: *Transdisciplinary Perspectives on Complex Systems: New Findings and Approaches*, Springer, 2017, pp. 85–113, http://dx.doi.org/10.1007/978-3-319-38756-7_4.
- [5] A. Rasheed, O. San, T. Kvamsdal, Digital twin: Values, challenges and enablers from a modeling perspective, *IEEE Access* 8 (2020) 21980–22012, <http://dx.doi.org/10.1109/ACCESS.2020.2970143>.
- [6] S.A. Niederer, M.S. Sacks, M. Girolami, K. Willcox, Scaling digital twins from the artisanal to the industrial, *Nat. Comput. Sci.* 1 (5) (2021) 313–320, <http://dx.doi.org/10.1038/s43588-021-00072-5>.
- [7] Y. Ye, Q. Yang, F. Yang, Y. Huo, S. Meng, Digital twin for the structural health management of reusable spacecraft: a case study, *Eng. Fract. Mech.* 234 (2020) 107076, <http://dx.doi.org/10.1016/j.engfracmech.2020.107076>.
- [8] M. Podskarbi, D.J. Knezevic, Digital twin for operations - present applications and future digital thread, in: *OTC Offshore Technology Conference*, Houston, Texas, 2020, <http://dx.doi.org/10.4043/30553-MS>.
- [9] A. Phua, C. Davies, G. Delaney, A digital twin hierarchy for metal additive manufacturing, *Comput. Ind.* 140 (2022) 103667, <http://dx.doi.org/10.1016/j.compind.2022.103667>.
- [10] M. Jans-Singh, K. Leeming, R. Choudhary, M. Girolami, Digital twin of an urban-integrated hydroponic farm, *Data-Cent. Eng.* 1 (2020) 20, <http://dx.doi.org/10.1017/dce.2020.21>.
- [11] A. Tzachor, S. Sabri, C.E. Richards, A. Rajabifard, M. Acuto, Potential and limitations of digital twins to achieve the sustainable development goals, *Nat. Sustain.* 5 (10) (2022) 822–829, <http://dx.doi.org/10.1038/s41893-022-00923-7>.
- [12] G. Arcieri, C. Hoelzl, O. Schwery, D. Straub, K.G. Papakonstantinou, E. Chatzi, Bridging POMDPs and Bayesian decision making for robust maintenance planning under model uncertainty: An application to railway systems, *Reliab. Eng. Syst. Saf.* (2023) 109496, <http://dx.doi.org/10.1016/j.res.2023.109496>.
- [13] M.G. Kapteyn, D.J. Knezevic, D. Huynh, M. Tran, K.E. Willcox, Data-driven physics-based digital twins via a library of component-based reduced-order models, *Internat. J. Numer. Methods Engrg.* 123 (13) (2020) 2986–3003, <http://dx.doi.org/10.1002/nme.6423>.
- [14] B. Moya, A. Badiás, I. Alfaro, F. Chinesta, E. Cueto, Digital twins that learn and correct themselves, *Internat. J. Numer. Methods Engrg.* 123 (13) (2022) 3034–3044, <http://dx.doi.org/10.1002/nme.6535>.
- [15] W. Sisson, P. Karve, S. Mahadevan, Digital twin approach for component health-informed rotorcraft flight parameter optimization, *AIAA J.* 60 (3) (2022) 1923–1936, <http://dx.doi.org/10.2514/1.J.060770>.
- [16] B. Moya, I. Alfaro, D. Gonzalez, F. Chinesta, E. Cueto, Physically sound, self-learning digital twins for sloshing fluids, *PLoS One* 15 (6) (2020) e0234569, <http://dx.doi.org/10.1371/journal.pone.0234569>.
- [17] K. Cherifi, V. Mehrmann, P. Schulze, Simulations in a digital twin of an electrical machine, 2022, <http://dx.doi.org/10.48550/arXiv.2207.02171>, arXiv preprint [arXiv:2207.02171](https://arxiv.org/abs/2207.02171).
- [18] M.G. Kapteyn, J.V. Pretorius, K.E. Willcox, A probabilistic graphical model foundation for enabling predictive digital twins at scale, *Nat. Comput. Sci.* 1 (5) (2021) 337–347, <http://dx.doi.org/10.1038/s43588-021-00069-0>.
- [19] D. Koller, N. Friedman, *Probabilistic Graphical Models: Principles and Techniques*, MIT Press, Cambridge, MA, 2009.
- [20] S.J. Russell, *Artificial Intelligence: A Modern Approach*, Pearson Education, London, UK, 2020.
- [21] M. Torzoni, A. Manzoni, S. Mariani, Structural health monitoring of civil structures: a diagnostic framework powered by deep metric learning, *Comput. Struct.* 271 (2022) 106858, <http://dx.doi.org/10.1016/j.compstruc.2022.106858>.
- [22] V. Giglioli, I. Venanzi, V. Poggioni, A. Milani, F. Ubertini, Autoencoders for unsupervised real-time bridge health assessment, *Comput.-Aided Civ. Infrastruct. Eng.* 38 (2023) 959–974, <http://dx.doi.org/10.1111/mice.12943>.
- [23] M. Torzoni, A. Manzoni, S. Mariani, A multi-fidelity surrogate model for structural health monitoring exploiting model order reduction and artificial neural networks, *Mech. Syst. Signal Process.* 197 (2023) 110376, <http://dx.doi.org/10.1016/j.ymsp.2023.110376>.
- [24] M. Torzoni, A. Manzoni, S. Mariani, A deep neural network, multi-fidelity surrogate model approach for Bayesian model updating in SHM, in: *European Workshop on Structural Health Monitoring*, Springer International Publishing, 2023, pp. 1076–1086, http://dx.doi.org/10.1007/978-3-031-07258-1_108.
- [25] L. Rosafalco, M. Torzoni, A. Manzoni, S. Mariani, A. Corigliano, A self-adaptive hybrid model/data-driven approach to SHM based on model order reduction and deep learning, in: *Structural Health Monitoring Based on Data Science Techniques*, Springer International Publishing, 2022, pp. 165–184, <http://dx.doi.org/10.1007/978-3-030-81716-9>.
- [26] S. Eftekhar Azam, S. Mariani, Online damage detection in structural systems via dynamic inverse analysis: A recursive Bayesian approach, *Eng. Struct.* 159 (2018) 28–45, <http://dx.doi.org/10.1016/j.engstruct.2017.12.031>.
- [27] A. Quarteroni, A. Manzoni, F. Negri, *Reduced Basis Methods for Partial Differential Equations: An Introduction*, Vol. 92, Springer, 2015, <http://dx.doi.org/10.1007/978-3-319-15431-2>.
- [28] M. Torzoni, M. Tezzele, Digital-twin-SHM, 2023, <https://github.com/mtezzele/digital-twin-SHM>.
- [29] G. Morcouc, Performance prediction of bridge deck systems using Markov chains, *J. Perform. Constr. Facil.* 20 (2) (2006) 146–155, [http://dx.doi.org/10.1061/\(ASCE\)0887-3828\(2006\)20:2\(146\)](http://dx.doi.org/10.1061/(ASCE)0887-3828(2006)20:2(146)).

- [30] P.D. Thompson, M.B. Johnson, Markovian bridge deterioration: developing models from historical data, *Struct. Infrastr. Eng.* 1 (1) (2005) 85–91, <http://dx.doi.org/10.1080/15732470412331289332>.
- [31] R.S. Sutton, A.G. Barto, *Reinforcement Learning: An Introduction*, MIT Press, Cambridge, MA, 2018.
- [32] T. Taddei, J.D. Penn, M. Yano, A.T. Patera, Simulation-based classification; a model-order-reduction approach for structural health monitoring, *Arch. Comput. Methods Eng.* 25 (2018) 23–45, <http://dx.doi.org/10.1007/s11831-016-9185-0>.
- [33] N.M. Newmark, A method of computation for structural dynamics, *J. Eng. Mech.* 85 (3) (1959) 67–94, <http://dx.doi.org/10.1061/JMCEA3.0000098>.
- [34] A. Teughels, J. Maec, G. De Roeck, Damage assessment by FE model updating using damage functions, *Comput. Struct.* 80 (25) (2002) 1869–1879, [http://dx.doi.org/10.1016/S0045-7949\(02\)00217-1](http://dx.doi.org/10.1016/S0045-7949(02)00217-1).
- [35] C. Farrar, K. Worden, *Structural Health Monitoring: A Machine Learning Perspective*, John Wiley & Sons, 2013, <http://dx.doi.org/10.1002/9781118443118>.
- [36] A. Pandey, M. Biswas, Damage detection in structures using changes in flexibility, *J. Sound Vib.* 169 (1) (1994) 3–17, <http://dx.doi.org/10.1006/jsvi.1994.1002>.
- [37] L. Rosafalco, A. Manzoni, S. Mariani, A. Corigliano, Fully convolutional networks for structural health monitoring through multivariate time series classification, *Adv. Model. Simul. Eng. Sci.* 7 (2020) 38, <http://dx.doi.org/10.1186/s40323-020-00174-1>.
- [38] L. Rosafalco, M. Torzoni, A. Manzoni, S. Mariani, A. Corigliano, Online structural health monitoring by model order reduction and deep learning algorithms, *Comput. Struct.* 255 (2021) 106604, <http://dx.doi.org/10.1016/j.compstruc.2021.106604>.
- [39] M. Torzoni, L. Rosafalco, A. Manzoni, S. Mariani, A. Corigliano, SHM under varying environmental conditions: an approach based on model order reduction and deep learning, *Comput. Struct.* 266 (2022) 106790, <http://dx.doi.org/10.1016/j.compstruc.2022.106790>.
- [40] F. Chinesta, A. Huerta, G. Rozza, K. Willcox, *Model reduction methods*, in: E. Stein, R. de Borst, T.J.R. Hughes (Eds.), *Encyclopedia of Computational Mechanics*, second ed., John Wiley & Sons, 2017, pp. 1–36.
- [41] P. Benner, S. Grivet-Talocia, A. Quarteroni, G. Rozza, W.H.A. Schilders, L.M. Silveira (Eds.), Volume 1: System- and Data-Driven Methods and Algorithms, De Gruyter, Berlin, Boston, 2021, <http://dx.doi.org/10.1515/9783110498967>.
- [42] G. Rozza, G. Stabile, F. Ballarin, in: G. Rozza, G. Stabile, F. Ballarin (Eds.), *Advanced Reduced Order Methods and Applications in Computational Fluid Dynamics*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2022, <http://dx.doi.org/10.1137/1.9781611977257>.
- [43] G. Kerschen, J.C. Golinval, Physical interpretation of the proper orthogonal modes using the singular value decomposition, *J. Sound Vib.* 249 (5) (2002) 849–865, <http://dx.doi.org/10.1006/jsvi.2001.3930>.
- [44] G. Kerschen, J.C. Golinval, A. Vakakis, L.A. Bergman, The method of proper orthogonal decomposition for dynamical characterization and order reduction of mechanical systems: An overview, *Nonlinear Dynam.* 41 (1) (2005) 147–169, <http://dx.doi.org/10.1007/s11071-005-2803-2>.
- [45] M. Guo, J.S. Hesthaven, Reduced order modeling for nonlinear structural analysis using Gaussian process regression, *Comput. Methods Appl. Mech. Engrg.* 341 (2018) 807–826, <http://dx.doi.org/10.1016/j.cma.2018.07.017>.
- [46] M. Tezzele, L. Fabris, M. Sidari, M. Sicchiero, G. Rozza, A multi-fidelity approach coupling parameter space reduction and non-intrusive POD with application to structural optimization of passenger ship hulls, *Internat. J. Numer. Methods Engrg.* 124 (5) (2023) 1193–1210, <http://dx.doi.org/10.1002/nme.7159>.
- [47] M. Sieber, C.O. Paschereit, K. Oberleithner, Spectral proper orthogonal decomposition, *J. Fluid Mech.* 792 (2016) 798–828, <http://dx.doi.org/10.1017/jfm.2016.103>.
- [48] A. Towne, O.T. Schmidt, T. Colonius, Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis, *J. Fluid Mech.* 847 (2018) 821–867, <http://dx.doi.org/10.1017/jfm.2018.283>.
- [49] A. Lario, R. Maulik, O.T. Schmidt, G. Rozza, G. Mengaldo, Neural-network learning of SPOD latent dynamics, *J. Comput. Phys.* 468 (2022) 111475, <http://dx.doi.org/10.1016/j.jcp.2022.111475>.
- [50] K.R. dos Santos, D.G. Giovanis, K. Kontolati, D. Loukrezis, M.D. Shields, Grassmannian diffusion maps based surrogate modeling via geometric harmonics, *Internat. J. Numer. Methods Engrg.* 123 (15) (2022) 3507–3529, <http://dx.doi.org/10.1002/nme.6977>.
- [51] C.M. Bishop, *Pattern Recognition and Machine Learning*, in: *Information Science and Statistics*, Springer, New York, NY, 2006.
- [52] I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning*, MIT Press, Cambridge, MA, 2016.
- [53] F. Negri, redbKIT, version 2.2, 2016, <http://redbkit.github.io/redbKIT>.
- [54] A. Ankan, A. Panda, pgmpy: Probabilistic graphical models using python, in: *Proceedings of the 14th Python in Science Conference (SCIPY 2015)*, Citeseer, 2015.
- [55] F. Chollet, et al., Keras, 2015, <https://keras.io>.
- [56] K. Papakonstantinou, M. Shinozuka, Planning structural inspection and maintenance policies via dynamic programming and Markov processes. Part II: POMDP implementation, *Reliab. Eng. Syst. Saf.* 130 (2014) 214–224, <http://dx.doi.org/10.1016/j.res.2014.04.006>.
- [57] M. Ülker-Kaustell, *Some Aspects of the Dynamic Soil-Structure Interaction of a Portal Frame Railway Bridge* (Ph.D. thesis), KTH Royal Institute of Technology, 2009.
- [58] T. Arvidsson, J. Li, *Dynamic Analysis of a Portal Frame Railway Bridge Using Frequency Dependent Soil Structure Interaction* (Ph.D. thesis), KTH Royal Institute of Technology, 2011.
- [59] European Committee for Standardization, Part 2: Traffic loads on bridges, in: *EN 1991-2 Eurocode 1: Actions on Structures*, 2003, pp. 66–74.
- [60] P. Battaglia, J.B.C. Hamrick, V. Bapst, A. Sanchez, V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner, C. Gulchre, F. Song, A. Ballard, J. Gilmer, G.E. Dahl, A. Vaswani, K. Allen, C. Nash, V.J. Langston, C. Dyer, N. Heess, D. Wierstra, P. Kohli, M. Botvinick, O. Vinyals, Y. Li, R. Pascanu, Relational inductive biases, deep learning, and graph networks, 2018, <http://dx.doi.org/10.48550/arXiv.1806.01261>, arXiv preprint arXiv:1806.01261.
- [61] X. Glorot, Y. Bengio, Understanding the difficulty of training deep feedforward neural networks, *J. Mach. Learn. Res.* 9 (2010) 249–256.
- [62] D. Kingma, J. Ba, Adam: A method for stochastic optimization, in: *3rd International Conference on Learning Representations*, Vol. 3, 2015, pp. 1–13.