



Research paper

Active digital twins via active inference

Matteo Torzoni ^a ,* , Domenico Maisto ^b , Andrea Manzoni ^c , Francesco Donnarumma ^b ,
Giovanni Pezzulo ^b , Alberto Corigliano ^a 

^a Department of Civil and Environmental Engineering, Politecnico di Milano, Milan, 20133, Italy^b Institute of Cognitive Sciences and Technologies, National Research Council, Rome, 00185, Italy^c MOX – Department of Mathematics, Politecnico di Milano, Milan, 20133, Italy

ARTICLE INFO

Keywords:

Digital twins

Active inference

Free energy principle

Structural health monitoring

ABSTRACT

Digital twins are transforming engineering and applied sciences by enabling real-time monitoring, simulation, and predictive analysis of physical systems and processes. However, conventional digital twins rely primarily on passive data assimilation, which limits their adaptability in uncertain and dynamic environments. This paper introduces the *active digital twin* paradigm, based on active inference. Active inference is a neuroscience-inspired Bayesian framework for probabilistic reasoning and predictive modeling that unifies inference, decision-making, and learning under a single free energy minimization objective. By modeling the dynamics of the coupled physical–digital system as a partially observable Markov decision process, active digital twins autonomously balance pragmatic exploitation (maximizing goal-directed utility) and epistemic exploration (actively resolving uncertainty). As action becomes an integral part of the inference process, active digital twins actively seek information to maintain synchronization with, and learn from their physical counterparts. The proposed framework is assessed through virtual experiments of structural health monitoring and predictive maintenance of a railway bridge. The application showcases the step-by-step construction of a generative model enabling bidirectional perception–action interaction. The results demonstrate that active digital twins exhibit superior exploration capabilities compared to traditional reactive approaches, enabling enhanced autonomy and resilience.

1. Introduction

Over the past decade, the digital twin (DT) paradigm has emerged as a transformative approach for monitoring, control, and decision support, enabling diagnostic and predictive capabilities that surpass those of traditional computational models. As outlined in [National Academy of Engineering and National Academies of Sciences, Engineering, and Medicine \(2024\)](#), DTs differ from both forward digital models and digital shadows ([Kritzinger et al., 2018](#)). The former are designed to simulate how input parameters and internal states influence system behavior to generate observable outputs, while the latter focus on data assimilation and model updating. A DT is a tailored virtual representation that captures key attributes of a physical system or process ([Ferrari and Willcox, 2024](#)). This digital representation dynamically synchronizes with its physical counterpart by continuously assimilating sensor data and providing predictive capabilities. Specifically, DTs enable the simulation of what-if scenarios, supporting predictive decision-making aimed at maximizing utility. This paper proposes active inference (AIF) ([Parr et al., 2022](#)) as a new paradigm for

DTs. By modeling the twin's evolution as a partially observable Markov decision process (POMDP) ([Russell, 2020](#)), the AIF agent achieves intelligent automation under the *free-energy principle* ([Friston, 2010](#); [Friston et al., 2017b](#)). This results in a unified mathematical framework for a new class of *active digital twins* (ADTs), equipped with spontaneous exploration capabilities.

Emerging from aeronautical and aerospace engineering ([Shafiq et al., 2020](#); [Digital Twin: Definition & Value — An AIAA and AIA Position Paper, 2020](#)), DT applications nowadays expand across several domains. These include structural health monitoring and predictive maintenance ([Torzoni et al., 2024](#); [Varetti et al., 2026](#); [Li et al., 2017](#)), additive manufacturing ([Phua et al., 2022](#)), smart cities ([Jans-Singh et al., 2020](#)), energy transition ([Reis et al., 2025](#)), urban sustainability ([Tzachor et al., 2022](#)), geotechnical engineering ([Cotoarbă et al., 2024](#)), subduction zone modeling ([Henneking et al., 2025](#)), railway infrastructure management ([Arcieri et al., 2024](#)), aerial vehicles monitoring and control ([McClellan et al., 2022](#); [Tezzele et al., 2024](#)), spacecraft operations in orbit ([Henao-Garcia et al., 2025](#)), personalized

* Corresponding author.

E-mail addresses: matteo.torzoni@polimi.it (M. Torzoni), domenico.maisto@istc.cnr.it (D. Maisto), andrea1.manzoni@polimi.it (A. Manzoni), francesco.donnarumma@istc.cnr.it (F. Donnarumma), giovanni.pezzulo@istc.cnr.it (G. Pezzulo), alberto.corigliano@polimi.it (A. Corigliano).

<https://doi.org/10.1016/j.engappai.2026.114519>

Received 17 June 2025; Received in revised form 19 February 2026; Accepted 14 March 2026

Available online 23 March 2026

0952-1976/© 2026 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Nomenclature

A	Observation model array
A^u	Observation model second modality O^u
A^u_{Epi}	Slice of A^u for the epistemic state
A^u_{Non-Epi}	Slice of A^u for the non-epistemic state
A^{Ωδ}	Observation model for first modality $O^{Ωδ}$
A^{Ωδ}_{Entropic}	Purely entropic version of $A^{Ωδ}$
A^{Ωδ}_{Epi}	Slice of $A^{Ωδ}$ for the epistemic state
A^{Ωδ}_{Non-Epi}	Slice of $A^{Ωδ}$ for the non-epistemic state
B	Transition model array
B^{Epi}	Transition model for D^{Epi} factor
B^δ	Transition model for $D^δ$ factor
B^Ω	Transition model for $D^Ω$ factor
ℬ	Categorical parameters over B
b	Dirichlet prior parameters over ℬ
\hat{b}	Variational approximation of b
\hat{b}^*	Optimized variational posterior of b
c	Prior preferences array
c^u	Prior preferences for second modality O^u
c^{Ωδ}	Prior preferences for first modality $O^{Ωδ}$
d	Initial state prior array
d^{Epi}	Initial state prior for D^{Epi} factor
d^δ	Initial state prior for $D^δ$ factor
d^Ω	Initial state prior for $D^Ω$ factor
D	Digital state space
D	Digital state
D^{Epi}	Epistemic switch digital state factor
D^δ	Damage magnitude digital state factor
D^Ω	Damage location digital state factor
d	Digital state realization
D_{KL}	Kullback–Leibler divergence
\mathbb{E}_Q	Expectation under variational posterior
Epi	Epistemic state
E	Young's modulus
F	Variational free energy
F^π	Variational free energy for policy π
F	Number of hidden state factors
G	Expected free energy across all policies
G^π	Expected free energy for policy π
\hat{G}^π	Ground-truth informed G^π
H	Entropy of variational posterior
L	Number of samples per vibration record
M	Number of observation modalities
Non – Epi	Non-epistemic state
N_{par}	Number of system-level parameters
N_s	Number of sensor locations
N_δ	Number of δ intervals
N_Ω	Number of damageable regions
O	Observation space
O	Predicted observations
O^{Exp}	Recorded observations
O^u	Second observation modality (action)
O^{Ωδ}	First observation modality (sensor data)
o	Predicted observations realization
o^{Exp}	Recorded observations realization
P	Number of feasible control policies
p	Probability distribution

\tilde{p}	Probability distribution biased by $\tilde{p}(O)$
$\tilde{p}(O)$	Prior preferences over observations
Q	Variational posterior
Q*	Optimized variational posterior
ℛ	Reward space
R	Reward
r	Reward realization
S	Physical state space
S	Physical state
s	Physical state realization
T	Time horizon
t	Time variable
t_c	Current time
t_p	Prediction time
U	Multivariate vibration recordings
u	Channel-wise vibration recordings
U	Action space
U	Control Action
U^{Epi}	Control factor for D^{Epi} factor
U^δ	Control factor for $D^δ$ factor
U^Ω	Control factor for $D^Ω$ factor
u	Control action realization
v	Train velocity
X	Generic tensor
X	Generic set
X	Generic random variable
Y	Generic random variable
y	Location of the stiffness reduction
α	Observational entropy coefficient
γ	Inverse temperature parameter
Δ^G	Expected free energy discrepancy metric
δ	Magnitude of the stiffness reduction
η	Learning rate
θ	Variational parameters
θ*	Optimized variational parameters
μ	System-level parameters
ν	Poisson's ratio
Π	Set of feasible control policies
π	Control policy
ρ	Material density
σ	Softmax function
φ	Generative model parameters
ψ	Axle-carried mass
Ω	Damageable subdomain
⊗	Outer product
ADT	Active digital twin
AIF	Active inference
Cat	Categorical distribution
CPT	Conditional probability table
DBN	Dynamic Bayesian network
Dir	Dirichlet distribution
DL	Deep learning
DN	Do nothing
DT	Digital twin
dof	Degree of freedom
EFE	Expected free energy
FOM	Full-order model
MA	Maintenance

PGM	Probabilistic graphical model
RE	Read sensors
RO	Restrict operations
ROM	Reduced-order model
SHM	Structural health monitoring
VFE	Variational free energy

medicine (Corral-Acero et al., 2020; Chaudhuri et al., 2023), and climate science (Bauer et al., 2021). Despite the growing interest in DTs, their implementation remains highly customized, typically tailored on the specific application, and often hard to deploy. The need for a widely accepted framework for DTs is therefore increasingly recognized in both research and industry. Kapteyn et al. (2021) proposed an application-agnostic formulation for describing coupled physical-digital systems that evolve dynamically over time and interact via observed data and control inputs. A key contribution of their work is the abstraction of the coupled dynamical system into a generalized representation, which serves as the foundation for a mathematical description of DTs. This abstraction is consistent with agent-based representations in POMDPs, typically formalized using probabilistic graphical models (Koller and Friedman, 2009).

We introduce ADTs with enhanced exploratory capabilities, employing AIF agents based on discrete generative models to leverage and significantly extend the abstraction of physical-digital systems by Kapteyn et al. (2021). Active inference is a theoretical framework integrating perception, decision-making, and learning within the unified objective of *free energy minimization* (Parr et al., 2022). An AIF agent maintains an internal generative model of its environment, continuously updating its beliefs in response to sensory inputs. By minimizing variational free energy, the agent simultaneously fulfills two objectives: reducing the divergence between predicted and preferred future observations, and resolving expected uncertainty about hidden states through action. This dual mechanism naturally balances *exploitation* of existing knowledge to achieve specific goals with *exploration*, i.e., the acquisition of new information. These two imperatives can be referred to using interchangeable terminology. The exploitative or pragmatic behavior is associated with terms such as goal-directed behavior or utility maximization, while the exploratory or epistemic behavior is described using terms such as information seeking, information gain, or uncertainty resolution. The AIF framework has been applied in diverse domains, from neuroscience (Friston et al., 2015; FitzGerald et al., 2015; Parr and Friston, 2027; Pezzulo et al., 2018; Van de Maele et al., 2024) — for modeling decision-making under uncertainty — to reinforcement learning (Fountas et al., 2020; Mazzaglia et al., 2021), collective behavior (Maisto et al., 2024; Heins et al., 2024), and robotics (Buckley et al., 2017; Lanillos et al., 2021; Taniguchi et al., 2023; Vijayaraghavan et al., 2025), demonstrating its versatility in modeling dynamic systems.

The generative model of an AIF agent functions as a self-updating engine that unifies the key aspects underpinning ADTs — namely, data assimilation, state estimation, prediction, planning, and learning — under a Bayesian framework that generalizes across applications. Furthermore, as demonstrated in the following sections, AIF agents naturally provide a mechanism for active information seeking, thereby unlocking the full potential of ADTs. When combined with goal-directed (pragmatic) behavior and possibly enhanced with learning capabilities, this information-seeking (epistemic) drive enables ADTs to engage in spontaneous exploration in response to (potentially critical) uncertainty, ultimately maximizing pragmatic utility.

Compared to alternative approaches for developing DTs, AIF offers remarkable advantages. Unlike reinforcement learning, which relies on trial-and-error exploration (often infeasible in real-world applications) and typically requires extensive datasets, AIF enables ADTs to infer

hidden states and optimize future behavior using a compact generative model. Moreover, the free energy minimization imperative of AIF balances information-seeking (epistemic) and goal-directed (pragmatic) behaviors, without the need for manually tuned reward functions or random exploration. These features make AIF particularly well-suited for adaptive and robust ADTs.

We present the ADT paradigm through an application in structural health monitoring and predictive maintenance of engineering structures. Given the potentially high life-cycle costs — economic, social, and safety — associated with such systems, adopting a DT perspective is crucial to enable condition-based or predictive maintenance practices, replacing traditionally employed time-based methods (Glaser and Tolman, 2008; Achenbach, 2009). To this end, non-destructive tests and in-situ inspections are inadequate for continuous and global monitoring. Conversely, by assimilating sensor data from permanent data collection systems, vibration-based structural health monitoring techniques enable automated damage identification and evolution tracking (Torzoni et al., 2022a; García-Macías and Ubertaini, 2022). This paradigm shift has the potential to unlock personalized monitoring, management, and maintenance programs (Thelen et al., 2022; Torzoni et al., 2025), offering numerous benefits throughout the system life-cycle — including more informed structural safety assessments, better resource allocation, and increased system availability (Torzoni, 2024).

A graphical abstraction of the computational flow is illustrated in Fig. 1. The end-to-end loop spans from the physical to the digital domain through data assimilation and inference, and then back to the asset through action and observation, while explicitly accounting for uncertainty quantification, propagation, and resolution. We refer to the monitored asset, whose physical state is hidden to the AIF agent and only indirectly accessible via the sensed structural response, as the *external generative process*. The asset state evolves over time according to physical laws influenced by both its internal properties and external factors. These external factors might encompass long-term degradation mechanisms caused by chemical, physical, or mechanical aging, as well as sudden changes, such as discrete damage events or maintenance interventions (Zakić et al., 1991).

The digital counterpart (AIF agent) is defined by an *internal generative model*, implemented as a probabilistic graphical model in the form of a dynamic Bayesian network (DBN) (Koller and Friedman, 2009; Russell, 2020). This factored representation provides a systematic way to maintain a posterior belief about latent variables that characterize the (hidden) structural health of the asset, such as damage presence, location, and severity, by continuously integrating new observations within a sequential Bayesian inference scheme. Belief updating is achieved by minimizing variational *free energy*, which measures the discrepancy between the model's predicted observations and the actual sensor data.

In parallel, the internal generative model supports the forward simulation of future states. This enables the ADT to evaluate “what-if” trajectories for structural health evolution, conditioned on its current beliefs. This forecasting step involves modeling not only the asset's physical dynamics but also the agent's control, represented as latent variables encoding sequences of future actions, usually termed policies (Friston et al., 2012). Policy selection is then framed as an optimization problem, where the agent seeks to minimize the *expected free energy* — a quantity that balances (i) selecting policies that align future observations with (pragmatic) goal-directed prior preferences, and (ii) resolving uncertainty about hidden states through (epistemic) information-seeking. This formalism unifies inference and control: posterior beliefs are updated via free energy minimization, while action sequences are selected to minimize expected free energy, converting the problem of decision-making into a problem of inference under the generative model. Once an action is executed, the generative process evolves, and the bidirectional perception-action cycle restarts.

Open questions: The limitations of conventional DTs, typically restricted to passive observation and open-loop simulation, have been recognized in Bounceur and Kara (2025). However, the challenge of

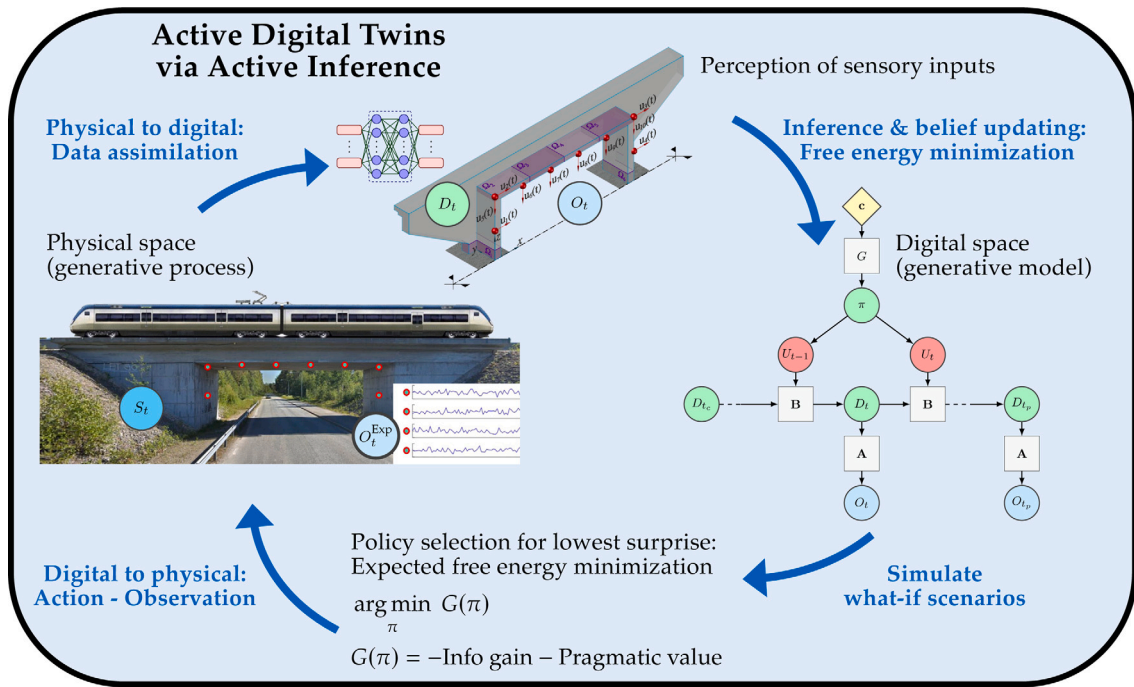


Fig. 1. Active digital twins via active inference — Graphical abstraction of the end-to-end information flow. The dichotomy between the external physical process generating observational data (i.e., the *generative process*) and the agent’s internal model (i.e., the *generative model*) is evident by the symmetry along the vertical axis. Meanwhile, the two forms of inference — digital state estimation and policy selection — exhibit a symmetry along the horizontal axis. A detailed schematic of generative models for both digital state and policy inference is presented in Fig. 4.

actively seeking information to enhance perception and learning remains largely unaddressed. Similar problems have long been studied in fields such as active vision, where perception is not limited to passively acquired images, but involves actively steering the sensing process to reduce uncertainty about the environment (Bajcsy, 1988; Aloimonos et al., 1988). Likewise, robotic systems dynamically adjust sensing devices to enhance environmental exploration (Thrun et al., 2005; Kaelbling et al., 1998), as in simultaneous localization and mapping (SLAM) tasks (Cadena et al., 2016). Similar perception mechanisms are embedded in autonomous driving systems, where sensor attention is dynamically allocated according to environmental conditions and contextual priorities (Grigorescu et al., 2020). Active digital twins build on the same principle of closing the loop between perception and action, enabling these systems to autonomously improve situational awareness, refine their internal models through active exploration, and proactively manage the environment evolution.

Novelty: While AIF has been widely applied in neuroscience, robotics, and decision theory, the novelty of this work lies in adopting it as a foundational framework for digital twins. By actively resolving uncertainty, ADTs foster adaptation and self-learning as means to maximize utility, in the spirit of intelligent automation (San et al., 2026). The balance between goal-directed (pragmatic) and information-seeking (epistemic) behaviors emerges naturally from the ADT’s generative model through expected free energy minimization (see also Fig. 1). The framework is formalized as a POMDP encoded through a probabilistic graphical model, providing an application-agnostic platform that extends previous abstractions of physical–digital systems (Kapteyn et al., 2021). We also provide a reproducible procedure for constructing the ADT’s generative model and demonstrate its advantages in structural health monitoring and predictive maintenance, highlighting behaviors unattainable by purely reactive digital twins.

The paper is organized as follows. Section 2 describes the POMDP encoding the coupled dynamics of the physical–digital system. Section 3 illustrates how AIF agents are used to realize ADTs. Section 4 assesses the proposed approach through simulated monitoring, management, and maintenance of a railway bridge, providing comparative

results for different AIF agents featuring increasingly rich behavior. The outcomes, along with current limitations and policy recommendations, are discussed in Section 5. Conclusions and future developments are finally outlined in Section 6.

2. Partially observable Markov decision process for digital twins

Fig. 2 illustrates the probabilistic graphical model — adapted from Kapteyn et al. (2021) — that represents the dynamic interaction between the physical and virtual domains. This abstraction is inspired by classical POMDP formulations (Russell, 2020). POMDPs are state–space models for decision-making in stochastic, partially observable environments, where system dynamics are typically described by Markov transition models. Unlike standard Markov decision processes, where a policy directly maps observable states to actions, POMDPs define the policy as a mapping from belief states — probabilistic representations of hidden states inferred from observations — to actions.

The graph in Fig. 2 is a DBN, in which circular nodes represent random variables, square nodes denote taken actions, and diamond-shaped nodes symbolize the objective function. All variables are defined at discrete time steps. Each time the DT is updated through the assimilation of new observational data, the DBN advances by one time step, with $t \in \{0, \dots, T\}$, where $t = 0$ marks the moment the DT enters operation, and $t = T$ defines its lifetime horizon. Nodes with bold outlines indicate observed quantities, while those with thin outlines correspond to latent variables that must be inferred. The DBN is sparsely connected, with edges encoding conditional dependencies among the variables. For an overview of the fundamentals of DBNs, the reader is referred to Koller and Friedman (2009) and Russell (2020).

Capital letters denote random variables associated with the quantities in our abstraction, the corresponding lowercase letters refer to their specific realizations, and subscripts indicate their time index. Calligraphic letters denote the set of possible values each quantity can assume. For instance, the hidden physical state is denoted as $S_t \sim p(s_t)$,

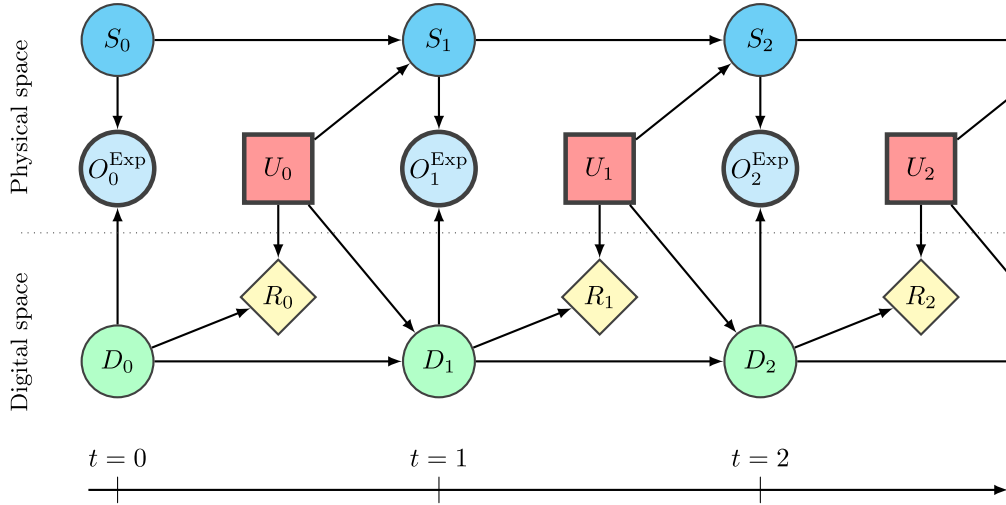


Fig. 2. Dynamic Bayesian network encoding the asset-twin dynamical system. Circular nodes represent random variables, square nodes denote taken actions, and diamond-shaped nodes symbolize the objective function. Nodes with bold outlines indicate observed quantities, while those with thin outlines represent latent variables to be inferred. Directed edges encode conditional dependencies between variables.

where s_t represents a particular realization at time t , and $p(s_t)$ defines the probability that $S_t = s_t$ for any possible state $s_t \in S$.

The digital state $D_t \sim p(d_t)$ is designed to capture the essential features of the (hidden) physical state that are relevant for diagnosis, prediction, and decision-making (Ferrari and Willcox, 2024). The digital state space D can represent a variety of information, including initial and/or boundary conditions, material properties, and other key characteristics to describe the asset under consideration.

The physical-to-digital information flow from S_t to D_t is mediated by the assimilation of observational data $O_t \sim p(o_t)$, enabling the inference of D_t . The observation space \mathcal{O} may include sensor measurements, inspection results, or diagnostic reports. Since the physical state S_t is only partially and indirectly observable, the digital state D_t encodes posterior beliefs over possible system configurations at time t , reflecting the evidence provided by the available observations (Kamariotis et al., 2024; Narouie et al., 2025). This perceptual process is realized through observation models — one for each observation modality — which relate digital states and observations in a probabilistic manner, such that O_t is modeled as stochastically generated from D_t . Throughout the paper, we will use both O_t and O_t^{Exp} to represent observations: O_t refers to predicted (expected) observations under the generative model, while O_t^{Exp} denotes actual sensor data. Belief updates are driven by minimizing the discrepancy between predicted and actual observations.

The updated digital state D_t informs the digital-to-physical information flow by guiding the selection of control actions to influence future physical states. In Fig. 2, $U_t \sim p(u_t)$ denotes a decision variable representing the action taken. The action space \mathcal{U} may include interventions that directly modify the physical state, adjustments to the operational conditions, or modifications to the observational process. Each action is associated with its own transition model — one for each digital state factor — across the digital state space, which serves as a control-dependent predictor that propagates the digital state beliefs forward in time.

Finally, the reward node $R_t \sim p(r_t)$ quantifies the performance of the asset-twin system within a reward space \mathcal{R} . These rewards assess the expected “quality” of DBN trajectories to guide action selection toward optimal outcomes. In general, reward values may represent real costs associated with states and actions, or abstract metrics tuned to steer the system toward the desired behavior.

Formally, a POMDP can be defined as a seven-tuple $\langle D, \mathcal{O}, \mathcal{U}, \mathcal{R}, \mathbf{A}, \mathbf{B}, \phi \rangle$, where: D denotes the space of beliefs over hidden states; \mathcal{O} is the space of possible observations; \mathcal{U} is the space of available actions; $\mathcal{R} : D \times \mathcal{U} \mapsto \mathbb{R}$ defines the reward function, which assigns a numerical

value to beliefs-action pairs; $\mathbf{A} : \mathcal{O} \times D \mapsto [0, 1]$ is the observation model, encoding the conditional observation likelihood $p(O_t | D_t; \phi)$, which represents beliefs about how hidden states give rise to observations; $\mathbf{B} : D \times D \times \mathcal{U} \mapsto [0, 1]$ is the transition model, encoding the conditional probability $p(D_t | D_{t-1}, U_{t-1}; \phi)$, which represents beliefs about the temporal evolution of hidden states conditioned on control actions; finally, ϕ is a vector of parameters of the POMDP model.

In the following, we assume that digital states, observational data, and control actions are defined over discrete and finite spaces. This implies that these variables can only take value on a finite set of discrete levels. Consequently, categorical distributions give a natural choice for representing the corresponding probability distributions. These latter assign a probability value between 0 and 1 to each discrete outcome, under the constraint that probabilities across all levels must sum to one, as they represent a complete and mutually exclusive set of realizations.

The joint probability distribution $p(O_t, D_t, U_t, R_t, \phi)$ over the POMDP factorizes — according to the chain rule of probability — into a product of categorical distributions (representing conditional likelihoods) and Dirichlet distributions (serving as priors). Numerically, these discrete distributions are organized as multidimensional arrays known as conditional probability tables (CPTs). The leading dimensions (rows) of a CPT correspond to the support of the random variable, while the lagging dimensions (columns) represent the conditioning variables. Each column specifies the probability distribution of a random variable given a particular configuration of its parent nodes, and the entries within each column sum to one, as they represent a complete set of mutually exclusive and exhaustive outcomes. If a node has no parents, its CPT reduces to a single column representing the prior probabilities of its possible values. The contents of these CPTs can be controlled through the model parameters included in ϕ .

The complete set of possible realizations of the unobserved variables — conditioned on observational data $O_{0:t_c}^{\text{Exp}} = o_{0:t_c}^{\text{Exp}}$ and control actions $U_{0:t_c} = u_{0:t_c}$ — from the initial time step $i = 0$ up to the current time t_c , with $t = 0, \dots, t_c$, can be extracted by leveraging the conditional independence assumptions implied by the graph structure in Fig. 2. The joint belief state can then be factorized according to the following sequential Bayesian inference formulation:

$$p(D_{0:t_c}, R_{0:t_c}, \phi | O_{0:t_c}^{\text{Exp}} = o_{0:t_c}^{\text{Exp}}, U_{0:t_c} = u_{0:t_c}) = p(\phi) p(D_0; \phi) \prod_{t=1}^{t_c} p(D_t | D_{t-1}, u_{t-1}; \phi) \prod_{t=0}^{t_c} p(O_t^{\text{Exp}} | D_t; \phi) p(R_t | D_t, u_t). \quad (1)$$

In Eq. (1), the term $p(\phi)$ represents the prior distribution over the model parameters ϕ ; inference over them typically evolves on a slower timescale than the inference of hidden states and control actions. $p(D_0; \phi)$ denotes the prior over the initial hidden states, representing the digital state belief at $t = 0$, before any observation is incorporated. The term $p(o_t^{\text{Exp}} | D_t; \phi)$ represents the sensory likelihood encoded in **A**. Similarly, $p(D_t | D_{t-1}, u_{t-1}; \phi)$ defines the transition likelihood encoded in **B**. Finally, $p(R_t | D_t, u_t)$ represents the likelihood of receiving a given reward, encapsulating the objective function evaluation. Note that selecting actions $U_t = u_t$ underpins solving the planning problem induced by the probabilistic graphical model. After forming a belief that measures the desirability of actions, such as $p(U_t | D_t)$, the actual action can be selected either as the best-point estimate or by sampling from this posterior, converting probabilistic control into a decision.

3. Active inference for digital twins

An attractive feature of AIF is that perception, learning, and action emerge as distinct manifestations of variational Bayesian inference (Parr et al., 2022). Perception, or state estimation, is accomplished through inference over dynamically evolving hidden states, conditioned on assimilated observations and past actions. Learning corresponds to the gradual inference of model parameters that capture the statistical regularities of the environment. Action, in turn, is realized by inferring a posterior distribution over policies and sampling actions accordingly.

In the following, we describe the use of AIF agents to “navigate” the POMDP underlying the DT problem, enabling the full potential of ADTs. Section 3.1 introduces the AIF generative model, which encodes the probabilistic assumptions about the underlying environment. Section 3.2 addresses digital state inference via variational free energy minimization. Section 3.3 covers policy inference and action selection through expected free energy minimization. Section 3.4 describes the slow-scale learning of the model parameters that define the AIF generative model. Finally, Section 3.5 discusses the active information-seeking (epistemic) behavior that characterizes ADTs.

3.1. Active inference generative model

In AIF, the set of probabilistic assumptions about how the environment (or *generative process*) produces observations (via the observation model **A**) and how actions influence the environment evolution (via the transition model **B**) is referred to as the POMDP generative model. This model is used to represent the joint distribution in Eq. (1), from current time t_c to a prediction horizon $t_p > t_c$. Specifically, for time- and space-discretized POMDPs, probabilistic estimates of future digital states and observations over the prediction time steps $t = t_c, \dots, t_p$ are computed as:

$$p(O_{t_c:t_p}, D_{t_c:t_p}, \phi | \pi) = p(\phi) p(D_{t_c}; \phi) \prod_{t=t_c+1}^{t_p} p(D_t | D_{t-1}, \pi; \phi) \prod_{t=t_c}^{t_p} p(O_t | D_t; \phi), \quad (2)$$

which reflects unrolling the AIF generative model of Fig. 3 over $t = t_c, \dots, t_p$. Compared to Eq. (1), the factorization in Eq. (2) introduces several modifications to align with the AIF framework. First, the control variable U is replaced by a policy π , defined as a sequence of control states $\pi = \{u_{t_c}, \dots, u_{t_p}\}$. The generative model in Eq. (2) is conditioned on a fixed policy π , which is how it is used for inference purposes. Policies are treated as latent variables to be inferred: the posterior over policies represents the agent beliefs about its intended actions, while single actions are realizations sampled from the posterior over control states. The policy-to-control mapping $p(U_t | \pi)$ assigns the control state at each time-step based on the selected policy.

The second modification concerns the omission of the reward variable R . In AIF, utility-maximization goals are encoded as a prior distribution $\tilde{p}(O_{t_c:t_p})$ over future observations. These preferences are

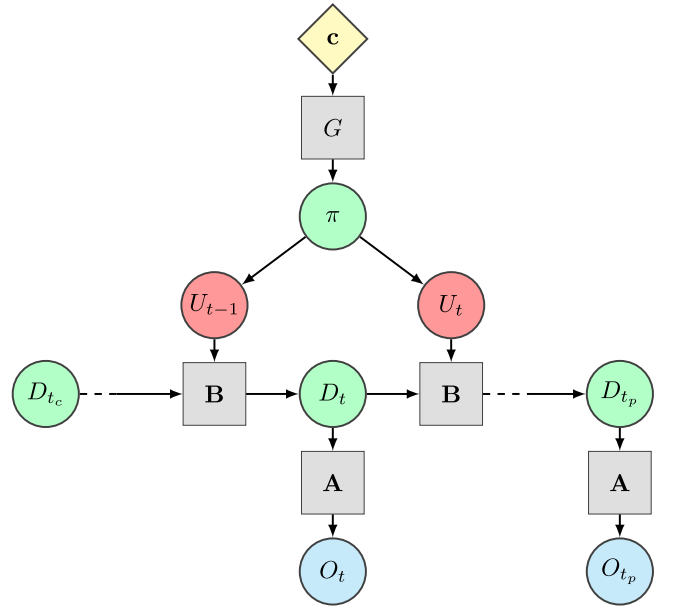


Fig. 3. Dynamic Bayesian network encoding the active inference generative model used to predict future digital states and observations under each policy. Circular nodes represent random variables, while the diamond-shaped node denotes prior preferences that reflect a goal-directed (pragmatic) objective. Gray square nodes represent parametrized operators of the generative model. Directed edges encode conditional dependencies between variables.

specified through an unconditional CPT $\mathbf{c} : \mathcal{O} \mapsto [0, 1]$. Such prior preferences guide policy selection toward goal-directed (pragmatic) behavior by favoring actions expected to produce preferred observations. This formal equivalence between rewards and priors eliminates the need for explicit cost functions. Further, it enables optimal control to be cast as an inference problem: the joint probability of observations, digital states, control states, and model parameters is maximized when the system samples from preferred observations. The square node G in the graph represents the expected free energy, which quantifies the desirability of each policy by incorporating both pragmatic and information-seeking (epistemic) components, as explained in Section 3.3. Finally, the initial prior $p(D_{t_c}; \phi)$ is typically represented by an unconditional CPT denoted as $\mathbf{d} : \mathcal{D} \mapsto [0, 1]$.

Graphically, the generative model illustrated in Fig. 3 shows several differences compared to the DBN in Fig. 2. This formulation focuses on predicting future digital states $D_{t_c:t_p}$ and sampling (or generating) sequences of potential observations $O_{t_c:t_p}$ based on the probabilistic structure encoded in **A** and **B**, conditioned on control actions $U_{t_c:t_p}$ that have not yet been executed. Accordingly, actions U_t are modeled as (circular) random variables rather than (square) decision nodes, since they represent what-if scenarios beyond data assimilation. Moreover, the same color is used to represent both digital states and control policies in the graph, as both are latent variables of the generative model. Equipped with this generative model — specified by the four-tuple $\langle \mathbf{A}, \mathbf{B}, \mathbf{c}, \mathbf{d} \rangle$ — AIF supports the inference over D_t , π , and ϕ , as described in the following sections.

3.2. Digital state inference via variational free energy minimization

Given an observation $O_{t_c}^{\text{Exp}} = o_{t_c}^{\text{Exp}}$, the underlying digital state D_{t_c} can be inferred by estimating a posterior distribution $p(D_{t_c} | O_{t_c}^{\text{Exp}} = o_{t_c}^{\text{Exp}})$, using Bayes' Rule:

$$p(D_{t_c} | O_{t_c}^{\text{Exp}} = o_{t_c}^{\text{Exp}}) = \frac{p(o_{t_c}^{\text{Exp}} | D_{t_c}) p(D_{t_c})}{p(o_{t_c}^{\text{Exp}})} = \frac{p(o_{t_c}^{\text{Exp}} | D_{t_c}) p(D_{t_c})}{\sum_{d_{t_c} \in \mathcal{D}} p(o_{t_c}^{\text{Exp}} | d_{t_c})}, \quad (3)$$

where the (generative model) joint distribution $p(o_{t_c}^{\text{Exp}}, D_{t_c})$ is factorized into a likelihood term $p(o_{t_c}^{\text{Exp}} | D_{t_c})$ and a prior $p(D_{t_c})$. The denominator $p(o_{t_c}^{\text{Exp}})$ is the marginal likelihood or model evidence, which captures the probability of observing $O_{t_c}^{\text{Exp}} = o_{t_c}^{\text{Exp}}$ under the generative model.

Since Bayesian inversion to estimate hidden states from observations is generally intractable, AIF employs variational inference (Murphy, 2023) as an approximate Bayesian method, trading exactness for computational tractability. Specifically, we define a tractable variational distribution $Q(D_{t_c}; \theta) : D \mapsto [0, 1]$, parametrized by θ , and optimize this surrogate distribution to make it as close as possible to the true posterior $p(D_{t_c} | O_{t_c}^{\text{Exp}} = o_{t_c}^{\text{Exp}})$. In our discrete POMDP setting, the variational parameters θ correspond to the relative frequencies of each category in the support of a random variable. This leads to the following optimization problem:

$$\theta^* = \arg \min_{\theta} D_{\text{KL}} \left[Q(D_{t_c}; \theta) \parallel p(D_{t_c} | o_{t_c}^{\text{Exp}}) \right], \quad (4)$$

where $D_{\text{KL}}[Q(X) \parallel P(X | Y)] = \mathbb{E}_Q[\ln Q(X) - \ln P(X | Y)]$ denotes the Kullback–Leibler (KL) divergence between the approximate posterior $Q(X)$ and the true posterior $P(X | Y)$, for two generic random variables X and Y . Here, \mathbb{E}_Q denotes the expectation with respect to the variational posterior. However, this objective remains intractable because it depends on the true posterior $p(D_{t_c} | o_{t_c}^{\text{Exp}})$ that we seek to approximate. To circumvent this, we reformulate the objective as the variational free energy (VFE):

$$\begin{aligned} F_{t_c}(\theta) &= D_{\text{KL}} \left[Q(D_{t_c}; \theta) \parallel p(D_{t_c} | o_{t_c}^{\text{Exp}}) \right] - \ln p(o_{t_c}^{\text{Exp}}) \\ &= \sum_D Q(D_{t_c}; \theta) \left[\ln \frac{Q(D_{t_c}; \theta)}{p(D_{t_c} | o_{t_c}^{\text{Exp}})} - \ln p(o_{t_c}^{\text{Exp}}) \right] \\ &= \mathbb{E}_Q \left[\ln Q(D_{t_c}; \theta) - \ln p(o_{t_c}^{\text{Exp}}, D_{t_c}) \right], \end{aligned} \quad (5)$$

which serves as an upper bound on the negative log marginal likelihood ($-\ln p(o_{t_c}^{\text{Exp}})$), also known as the Bayesian surprise. Minimizing VFE thus brings the variational posterior closer to the true posterior while simultaneously increasing the marginal likelihood of the observation. The VFE objective leads to the following final form of the optimization problem:

$$\theta^* = \arg \min_{\theta} F_{t_c}(\theta). \quad (6)$$

At convergence, if $Q^*(D_{t_c}; \theta^*)$ exactly matches the true posterior, the KL divergence vanishes, i.e., $D_{\text{KL}} = 0$, and the VFE equals surprise: $F_{t_c} = -\ln p(o_{t_c}^{\text{Exp}})$. By further minimizing surprise, the VFE then provides a useful objective not only for inference but also for learning the parameters of the generative model. The underlying rationale is that AIF agents aim to avoid surprising observations, and minimizing surprise is equivalent to maximizing model evidence.

With reference to the generative model formulation (2), instantaneous inference over digital states involves approximating the true posterior $p(D_{t_c} | O_{t_c}^{\text{Exp}} = o_{t_c}^{\text{Exp}}, D_{t_{c-1}}, U_{t_{c-1}} = u_{t_{c-1}})$. This inference is conditioned on the current observation $O_{t_c}^{\text{Exp}} = o_{t_c}^{\text{Exp}}$, the previous (posterior) distribution over digital states $D_{t_{c-1}}$, and the previously executed action $U_{t_{c-1}} = u_{t_{c-1}}$, as:

$$\begin{aligned} \theta^* &= \arg \min_{\theta} F_{t_c}(\theta) \\ &= \arg \min_{\theta} \mathbb{E}_Q \left[\ln Q(D_{t_c}; \theta) - \ln p(o_{t_c}^{\text{Exp}}, D_{t_c} | D_{t_{c-1}}, u_{t_{c-1}}; \phi) \right] \\ &= \arg \min_{\theta} \mathbb{E}_Q \left[\ln Q(D_{t_c}; \theta) \right. \\ &\quad \left. - \ln \left(p(o_{t_c}^{\text{Exp}} | D_{t_c}; \phi) p(D_{t_c} | D_{t_{c-1}}, u_{t_{c-1}}; \phi) \right) \right]. \end{aligned} \quad (7)$$

The optimization problem in Eq. (7) is solved using fixed-point iteration (Wainwright and Jordan, 2008), under the assumption of temporal factorization, where variational posteriors at different time steps are conditionally independent. As a result, the full VFE across

trajectories decomposes into a sum of single-time-step free energies, enabling independent optimization at each time point.

The posterior $Q(D_t; \theta)$ at a given time step t can be further factorized across F independent hidden state factors $D = \{D^1, \dots, D^F\}$, following the mean-field approximation (Bishop, 2006):

$$Q(D_t; \theta) = \prod_{f=1}^F Q(D_t^f; \theta), \quad (8)$$

where $Q(D_t^f; \theta)$ denotes the posterior over the f th hidden state factor, $f = 1, \dots, F$. These factors may represent distinct aspects of the generative process, potentially varying in dimensionality, transition dynamics, and association with specific observation modalities. Similarly, observations can be structured into M distinct modalities $O = \{O^1, \dots, O^M\}$, where each O^m , $m = 1, \dots, M$, corresponds to a separate sensory channel used by the agent at each time step. For example, in a DT application for the human health, one hidden state factor may represent a patient's metabolic state, while another factor could encode cardiovascular function. Correspondingly, observation modalities may include blood glucose readings and heart rate measurements, each providing information about different latent physiological processes.

In this multi-modal, multi-factor setup, the observation likelihood array \mathbf{A} becomes a collection of M sub-arrays $\mathbf{A} = \{\mathbf{A}^1, \dots, \mathbf{A}^M\}$, with each \mathbf{A}^m , $m = 1, \dots, M$, representing the observation model for the m th modality. Each sub-array encodes the likelihood $p(O^m | D^1, \dots, D^F; \phi)$, capturing the dependency of that observation modality on the hidden state factors. Similarly, the transition model \mathbf{B} is represented as a collection of F sub-arrays $\mathbf{B} = \{\mathbf{B}^1, \dots, \mathbf{B}^F\}$, under the assumption that hidden state factors evolve independently without influencing each other. Each \mathbf{B}^f , $f = 1, \dots, F$, encodes the dynamics $p(D_t^f | D_{t-1}^f, u_{t-1}^f; \phi)$, conditioned on the previous state and action for that factor. Note that control states are factorized analogously to hidden states, such that $U = \{U^1, \dots, U^F\}$. Each control factor U^f governs the transitions of the corresponding digital state factor D^f , with a dimensionality matching the number of possible control actions applicable to that aspect of the system.

This factored structure enables the encoding of complex conditional dependencies while significantly reducing memory requirements. For instance, if the model employs two separate hidden state factors to represent the location and identity of a phenomenon, the memory requirements for the factored representation scale linearly with the dimensionality of the two factors. An additional advantage of this factorization lies in its interpretability: by explicitly designing digital state factors to reflect intuitive features of the environment, the resulting generative model becomes more transparent and modular. In contrast, explicitly enumerating all possible combinations of “where” and “what” would incur polynomial memory complexity.

The marginal variational posteriors for each hidden state factor at the current time t_c are computed via mean-field fixed-point iteration (Wainwright and Jordan, 2008). The algorithm proceeds by setting the gradient of the VFE $F_{t_c}(\theta)$ to zero, and iteratively solving for each factorized component $Q(D_{t_c}^f; \theta)$, for $f = 1, \dots, F$. A detailed derivation of this procedure can be found in Heins et al. (2022).

Note that in the AIF framework, there is no need to introduce an explicit node for representing quantities of interest, unlike the abstraction of physical–digital systems proposed in Kapteyn et al. (2021). In their probabilistic graphical model, these variables are represented by a dedicated node and predicted from the updated digital state via the computational models comprising the DT. In contrast, under the AIF framework, such a node is redundant, as quantities of interest are naturally embedded within the observational data node. When observational evidence is unavailable for a particular modality, inference simply remains uninformed in that dimension of the observation space. Nevertheless, the updated digital state can still be used to predict expected values across any observation channel — whether observed or unobserved — via the corresponding observation model. The models may, in principle, incorporate arbitrarily complex forward mappings, ranging from high-fidelity physics-based simulators to purely data-driven surrogates or hybrid combinations of the two.

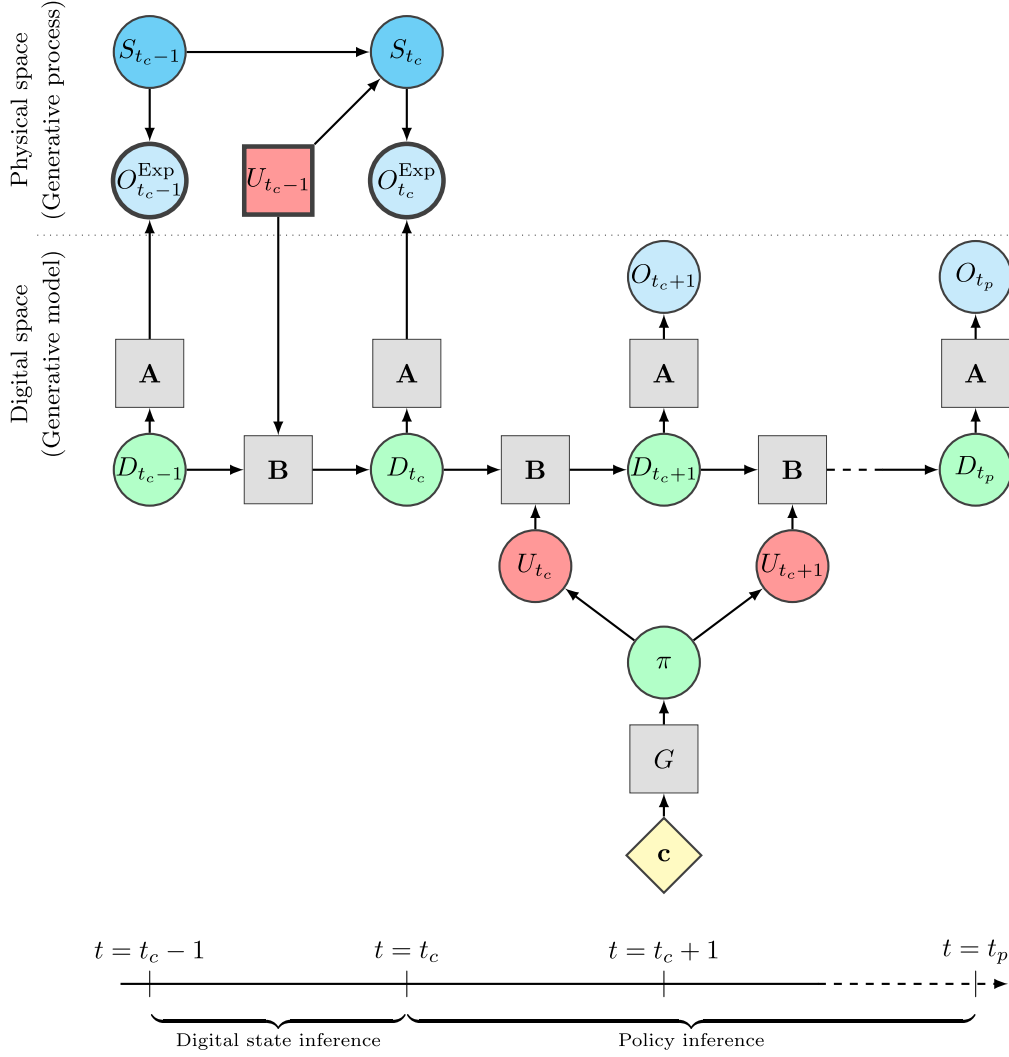


Fig. 4. A dynamic Bayesian network illustrating the use of active inference generative models to navigate the partially observable Markov decision process underlying the digital twin problem. Circular nodes represent random variables, red square nodes denote taken actions, gray square nodes represent parametrized operators of the generative model, and the diamond-shaped node symbolizes prior preferences that reflect a goal-directed (pragmatic) objective. Nodes with bold outlines indicate observed quantities, while those with thin outlines represent latent variables to be inferred. Directed edges encode conditional dependencies between variables. The upper left-to-right path represents the evolution of the physical space, while the lower path depicts the evolution of the digital space. Digital state inference is performed at current time t_c , whereas policy inference involves propagating the updated digital state from t_c to the prediction time t_p .

3.3. Policy inference-action selection via expected free energy minimization

Given the updated variational posterior over the digital state $Q^*(D_{t_c}; \theta^*)$, policy inference involves evaluating the quality of each admissible policy comprising future actions over a prediction horizon $t = t_c, \dots, t_p$. In AIF, the desirability of (or preference for) each policy is quantified through the expected free energy (EFE). The EFE is the central quantity driving the behavior of ADTs and is formulated to evaluate sequences of actions (or policies) both on goal-directed (pragmatic) and information-seeking (epistemic) behaviors. Like the VFE, the EFE is a function of observations, hidden states, and policies. However, different from the VFE, it pertains to sequences of *future* actions, where no actual observations are yet available, and it includes expectations over future digital states and future observations generated by the generative model.

The use of AIF generative models for digital state inference and policy inference is graphically summarized in Fig. 4. Digital state inference integrates the prior belief at time $t_c - 1$ with the observational

data assimilated at t_c . In contrast, policy inference entails predictive modeling over the horizon $t = t_c, \dots, t_p$, where the generative model operates without access to future sensory data or executed actions from the interfacing generative process.

The EFE associated to a generic policy π is defined as:

$$G^\pi = \mathbb{E}_{Q(O_{t_c:t_p}, D_{t_c:t_p} | \pi)} \left[\ln Q(D_{t_c:t_p} | \pi) - \ln \tilde{p}(O_{t_c:t_p}, D_{t_c:t_p} | \pi) \right], \quad (9)$$

where, for simplicity, we omit the explicit dependence of the variational posterior on the variational parameters θ , denoting it simply as $Q(D_t)$. Similarly, we omit the dependency of the generative model on the model parameters ϕ . In Eq. (9), $\tilde{p}(O_t, D_t | \pi) = p(D_t | O_t, \pi) \tilde{p}(O_t)$ defines a generative model biased by the predictive prior over observations $\tilde{p}(O_t)$. This construction integrates the prior preferences encoded in c into the inference process (described below), enabling the AIF agent to act in ways that maximize the likelihood of preferred outcomes.

Given the assumed conditional independence of variational posteriors across time, the EFE at a generic time step $t \in \{t_c, \dots, t_p\}$ for policy

π is given by:

$$\begin{aligned}
 G_t^\pi &= \mathbb{E}_{Q(O_t, D_t | \pi)} [\ln Q(D_t | \pi) - \ln \tilde{p}(O_t, D_t | \pi)] \\
 &= \underbrace{-\mathbb{E}_{Q(O_t | \pi)} [\mathbb{D}_{\text{KL}} [Q(D_t | O_t, \pi) \| Q(D_t | \pi)]]}_{\text{Epistemic value (information gain)}} - \underbrace{\mathbb{E}_{Q(O_t | \pi)} [\ln \tilde{p}(O_t)]}_{\text{Pragmatic value (utility)}} \\
 &\quad + \underbrace{\mathbb{E}_{Q(O_t | \pi)} [\mathbb{D}_{\text{KL}} [Q(D_t | O_t, \pi) \| p(D_t | O_t, \pi)]]}_{\text{Expected variational approximation error } (\geq 0)},
 \end{aligned} \tag{10}$$

with the complete derivation provided in the Appendix, as adapted to the ADT framework from Heins et al. (2022). In Eq. (10), the first term denotes the epistemic value (Friston et al., 2015), which promotes information-seeking behavior. It favors policies under which the agent is expected to *explore* states that yield high information gain about the digital state. This gain is quantified as the divergence between predicted digital states conditioned and unconditioned on observations under the same policy. The second term corresponds to the pragmatic value, which reflects goal-directed behavior. It favors policies that lead the agent to states expected to generate outcomes aligned with prior preferences $\tilde{p}(O_t)$. The final term captures the expected approximation error — the divergence between the true digital state posterior and its variational approximation — which is typically assumed to be negligible.

The epistemic drive in Eq. (10) is a crucial component that enables ADTs to exhibit spontaneous exploratory behavior. Epistemic actions in ADTs encompass decisions that gather information or improve the digital state observability. These may include, for instance, installing new sensors, scheduling targeted inspections, or testing model predictions. For example, in a manufacturing ADT, the agent might deliberately vary process parameters within safe limits to resolve uncertainty about machine wear dynamics. In a personalized medicine context, the ADT might recommend a low-risk diagnostic test to disambiguate between competing hypotheses about a patient's physiological condition. In both cases, the primary objective of these actions is not immediate (pragmatic) utility maximization, but rather to refine the generative model and enhance the understanding of the environment.

The EFE of temporally deep policies is given by the sum of time step-specific contributions:

$$G^\pi = \sum_{t=t_c}^{t_p} G_t^\pi, \tag{11}$$

where each term is evaluated based on the agent's predictive beliefs over future digital states and observations. The computation begins from the current posterior belief $Q^*(D_{t_c})$, which is then propagated over the prediction horizon $t = t_c, \dots, t_p$ using the policy-specific transition and observation models. This process generates the posterior predictive densities $Q(O_{t_c:t_p}, D_{t_c:t_p} | \pi)$, which are subsequently used to evaluate the goal-directed (pragmatic) and information-seeking (epistemic) values at each time step.

Let $\Pi = \{\pi_1, \dots, \pi_P\}$ denote the set of P feasible policies, constructed through the combinatorial enumeration of sequences of actions from the action space \mathcal{U} over the time horizon $t = t_c, \dots, t_p$. The EFE vector $\mathbf{G} = (G^{\pi_1}, \dots, G^{\pi_P})^\top \in \mathbb{R}^P$, which assigns a scalar EFE to each policy, defines a prior over policies according to:

$$p(\pi) = \sigma(-\gamma \mathbf{G}), \tag{12}$$

where $\sigma(x) = \frac{\exp(x)}{\sum_x \exp(x)}$ is the Softmax function, and $\gamma \in \mathbb{R}^+$ is an inverse temperature parameter that modulates the precision over policies. Higher γ values yield more deterministic preferences.

Under the prior (12), AIF agents perform policy inference by optimizing a variational posterior over policies $Q(\pi)$ (Heins et al., 2022), to minimize the following VFE expansion over the prediction horizon

$$t = t_c, \dots, t_p:$$

$$\begin{aligned}
 F_{t_c:t_p} &= \mathbb{E}_{Q(D_{t_c:t_p}, \pi)} [\ln Q(D_{t_c:t_p}, \pi) - \ln p(O_{t_c:t_p}, D_{t_c:t_p}, \pi)] \\
 &= \mathbb{E}_{Q(D_{t_c:t_p}, \pi)} [\ln Q(D_{t_c:t_p} | \pi) + \ln Q(\pi) - \ln p(O_{t_c:t_p}, D_{t_c:t_p} | \pi) - \ln p(\pi)] \\
 &= \mathbb{E}_{Q(\pi)} [\ln Q(\pi) - \ln p(\pi)] \\
 &\quad + \mathbb{E}_{Q(\pi)} [\mathbb{E}_{Q(D_{t_c:t_p} | \pi)} [\ln Q(D_{t_c:t_p} | \pi) - \ln p(O_{t_c:t_p}, D_{t_c:t_p} | \pi)]] \\
 &= \mathbb{D}_{\text{KL}} [Q(\pi) \| p(\pi)] + \mathbb{E}_{Q(\pi)} [\mathcal{F}_{t_c:t_p}^\pi],
 \end{aligned} \tag{13}$$

which measures the KL divergence between the approximate posterior $Q(D_{t_c:t_p}, \pi)$ and the generative model $p(O_{t_c:t_p}, D_{t_c:t_p}, \pi)$ as a sum of two contributions. The first is the KL divergence between the variational posterior over policies and the corresponding prior (12), thereby incorporating the EFE into the inference process. The second term is a policy-weighted average of the free energy across all policies, where $\mathcal{F}_{t_c:t_p}^\pi$ denotes the free energy associated with a single policy π :

$$\begin{aligned}
 \mathcal{F}_{t_c:t_p}^\pi &= -\mathbb{E}_{Q(D_{t_c:t_p} | \pi)} [\ln p(O_{t_c:t_p}, D_{t_c:t_p} | \pi) - \ln Q(D_{t_c:t_p} | \pi)] \\
 &= -\mathbb{E}_{Q(D_{t_c:t_p} | \pi)} [\ln p(O_{t_c:t_p}, D_{t_c:t_p} | \pi)] - H [Q(D_{t_c:t_p} | \pi)],
 \end{aligned} \tag{14}$$

with $H [Q(D_{t_c:t_p} | \pi)] = \mathbb{E}_{Q(D_{t_c:t_p} | \pi)} [-\ln Q(D_{t_c:t_p} | \pi)]$ being the variational posterior entropy, which quantifies the uncertainty in the beliefs about future digital states under policy π .

By evaluating each policy independently and computing its associated free energy, the optimal posterior $Q^*(\pi)$ is obtained by minimizing the total VFE $F_{t_c:t_p}$ with respect to $Q(\pi)$. This is achieved by enforcing the stationarity of $F_{t_c:t_p}$ with respect to $Q(\pi)$, leading to a Softmax distribution through the following update rule:

$$Q^*(\pi) = \arg \min_{Q(\pi)} F_{t_c:t_p} = \sigma(\ln p(\pi) - \mathcal{F}_{t_c:t_p}^\pi), \tag{15}$$

assigning higher probability to policies with lower free energy while remaining close to the prior.

The posterior over control states $Q^*(U_t)$ is formed by marginalizing over policies as follows:

$$Q^*(U_t) = \sum_{\pi \in \Pi} p(U_t | \pi) Q^*(\pi), \tag{16}$$

where $p(U_t | \pi)$ defines a deterministic mapping from policies to control states. The actual action $U_{t_c} = u_{t_c}$ to be executed on the system can eventually be selected either as the maximum a-posteriori estimate or by sampling from $Q^*(U_{t_c})$.

3.4. Learning of the generative model via parameter inference

In this section, we describe the learning of the parameters ϕ that define the AIF generative model, based on the outcomes of inference. "Learning" ϕ is a generative model's parameter updating occurring at a slower timescale than the faster inference processes for digital states and policies. Nevertheless, the update equations for ϕ follow the same variational principles of digital state inference, where a variational posterior over ϕ is optimized through VFE minimization.

In our discrete setting, posterior inference over ϕ is performed by parametrizing the likelihood and prior distributions of the generative model with Dirichlet distributions, following an approach similar to Tezzele et al. (2024) and Varetti et al. (2026). The choice to treat model parameters ϕ as the parameters of Dirichlet distributions is motivated by their conjugacy to the categorical distribution. This formulation enables online learning via closed-form Bayesian updates, allowing evidence about the system response to actions to be incorporated efficiently, while ensuring that the posterior remains within the Dirichlet family. The approach is computationally scalable and supports continual refinement of ADTs, even when initialized with potentially inaccurate or uncertain priors and likelihoods.

By decomposing ϕ into subsets corresponding to the categorical and Dirichlet parameters associated with the arrays \mathbf{A} , \mathbf{B} , and \mathbf{d} , we can highlight the stochastic parametrization underlying each likelihood and prior distribution. In this work, we focus on learning the parameters of the \mathbf{B} array, although the same reasoning can be applied identically to \mathbf{A} and \mathbf{d} . The stochastic parametrization of the transition model $\mathbf{B} \in \mathbb{R}^{|\mathcal{D}| \times |\mathcal{D}| \times |\mathcal{U}|}$, encoding the transition distribution $p(D_t | D_{t-1}, u_{t-1}; \mathfrak{B})$, is given by the tensor of categorical parameters $\mathfrak{B} \in \mathbb{R}^{|\mathcal{D}| \times |\mathcal{D}| \times |\mathcal{U}|}$, such that

$$D_t | D_{t-1}, u_{t-1}; \mathfrak{B} \sim \text{Cat}(\mathfrak{B}), \quad (17)$$

$$p(\mathfrak{B}) = \prod_{d \in \mathcal{D}} \prod_{u \in \mathcal{U}} p(\mathfrak{B}_{\cdot, d, u}), \quad \mathfrak{B}_{\cdot, d, u} \sim \text{Dir}(\mathbf{b}_{\cdot, d, u}), \quad (18)$$

where $|\mathcal{X}| \in \mathbb{N}$ denotes the cardinality of a generic set \mathcal{X} ; $\text{Cat}(\cdot)$ and $\text{Dir}(\cdot)$ denote categorical and Dirichlet distributions, respectively; the notation $\mathbf{X}_{\cdot, j, k}$ refers to the j th column of the k th slice of a generic tensor \mathbf{X} . The tensor $\mathbf{b} \in \mathbb{R}^{|\mathcal{D}| \times |\mathcal{D}| \times |\mathcal{U}|}$ collects the (positive) concentration parameters defining the Dirichlet prior over \mathfrak{B} . These parameters can be interpreted as pseudo-counts representing prior beliefs about the frequency of each state transition given a digital state and an action. For notational simplicity, we assume the generative model is not factorized into multiple digital state factors or observation modalities. However, the formulation extends naturally to multi-modal, multi-factor settings by introducing additional parametrized dimensions.

By specializing ϕ as the parametrization of \mathbf{B} , the generative model (2) becomes:

$$p(O_{t_c:t_p}, D_{t_c:t_p}, \mathfrak{B}, \pi) = p(\mathfrak{B}) p(\pi) p(D_{t_c}) \prod_{t=t_c+1}^{t_p} p(D_t | D_{t-1}, \pi; \mathfrak{B}) \prod_{t=t_c}^{t_p} p(O_t | D_t). \quad (19)$$

Learning is formulated as the approximate inference of \mathfrak{B} by minimizing the VFE with respect to its approximate posterior $Q(\mathfrak{B})$. The full variational posterior is assumed to factorize as:

$$Q(D_{t_c:t_p}, \mathfrak{B}, \pi) = Q(\mathfrak{B}) Q(\pi) \prod_{t=t_c}^{t_p} Q(D_t | \mathfrak{B}, \pi), \quad (20)$$

where the variational distribution $Q(\mathfrak{B})$ is modeled as

$$Q(\mathfrak{B}) = \prod_{d \in \mathcal{D}} \prod_{u \in \mathcal{U}} Q(\mathfrak{B}_{\cdot, d, u}), \quad Q(\mathfrak{B}_{\cdot, d, u}) = \text{Dir}(\hat{\mathbf{b}}_{\cdot, d, u}), \quad (21)$$

and $\hat{\mathbf{b}} \in \mathbb{R}^{|\mathcal{D}| \times |\mathcal{D}| \times |\mathcal{U}|}$ plays the same role as \mathbf{b} in parametrizing the Dirichlet distribution, while being treated as variational parameters to be optimized. Accordingly, the full VFE objective for the generative model (19) is given by:

$$\mathcal{F}_{t_c:t_p} = \mathbb{E}_{Q(D_{t_c:t_p}, \mathfrak{B}, \pi)} \left[\ln Q(D_{t_c:t_p}, \mathfrak{B}, \pi) - \ln p(O_{t_c:t_p}, D_{t_c:t_p}, \mathfrak{B}, \pi) \right], \quad (22)$$

which decomposes into sums of KL and expected log-likelihood terms following the factorization induced by the generative model and the variational posterior.

The update rule for the Dirichlet parameters \mathbf{b} directly follows from conjugacy. Specifically, given the Dirichlet prior parameters \mathbf{b} over the generative model, the digital state posterior $Q^*(D_{t_c})$ at the current time step t_c , the previous digital state posterior $Q^*(D_{t_c-1})$, and the action $U_{t_c-1} = u_{t_c-1}$ taken at the previous time step, the fixed-point update rule for the variational posterior Dirichlet parameters $\hat{\mathbf{b}}$ is:

$$\hat{\mathbf{b}}_{\cdot, \cdot, u_{t_c-1}}^* = \mathbf{b}_{\cdot, \cdot, u_{t_c-1}} + \eta (Q^*(D_{t_c}) \otimes Q^*(D_{t_c-1})), \quad (23)$$

which corresponds to an update applied to the u_{t_c-1} th slice of $\hat{\mathbf{b}}$, where \otimes denotes the outer product, and $\eta \in \mathbb{R}$, with $0 \leq \eta \leq 1$, is a learning rate parameter that scales the update step.

3.5. Epistemic behavior of active digital twins

If the agent also maintains a variational posterior over the model parameters $Q(\phi)$, as discussed in Section 3.4, the EFE expression (10) can be extended to capture the epistemic value associated with expected information gain not only over digital states but also over ϕ . The full derivation is provided in the Appendix, and it reveals an additional epistemic term that quantifies the expected information gain about the parameters governing the categorical distributions associated with the \mathbf{A} , \mathbf{B} , and \mathbf{d} arrays. When the AIF agent maintains and updates beliefs over these parameters, this term steers policy inference toward action-observation trajectories that are expected to produce informative updates to the generative model. We point out that the epistemic value over ϕ is not exploited in the numerical demonstrations presented in Section 4. Nevertheless, it is retained in the formulation of the ADT framework, as this capability may enable essential functionalities depending on the context and specific application objectives.

Epistemic actions aimed at refining the generative model can be regarded as forms of autonomous calibration, wherein the ADT steers its operation into underexplored regimes or perturbs its environment to test and improve its generative model. For instance, a sensor might be temporarily activated solely to evaluate its reliability while updating a likelihood model deemed uncertain. This behavior highlights the distinction between passive and active learning: while passive learning consists in assimilating externally provided or randomly encountered data, active learning reflects the strategic initiation of data acquisition to accelerate model refinement and enhance future decision-making. The information-seeking (epistemic) behavior of ADTs thus emerges from their capacity for self-adaptive inference and learning, pursued alongside goal-directed (pragmatic) objectives. This dual optimization is embedded in policy inference through EFE minimization, which unifies goal-directed exploration and utility maximization within a single computational framework that moves beyond the passive replication of physical systems.

3.6. Algorithmic description

An algorithmic description of a single step of the AIF loop for ADTs is provided in Algorithm 1. Given the generative model, an observation sampled from the generative process, the posterior over digital states from the previous time step, and the action taken at the previous time step, one step of the loop involves: (1) performing inference over digital states based on the new observation; (2) using the posterior belief over digital states to perform policy inference and select the next action; (3) updating the generative model through learning informed by inference results.

4. Numerical demonstrations

This section demonstrates the proposed methodology through the simulated monitoring, management, and maintenance planning of the Hörnefors railway bridge (Ülker-Kaustell, 2009). Although this case study focuses specifically on structural health monitoring (SHM), the underlying framework is similarly applicable to a wide range of systems or domains.

Section 4.1 introduces the monitored physical asset. Section 4.2 describes the composition of the handled vibration data and the numerical models used to generate labeled examples under various damage scenarios. Section 4.3 outlines the assimilation of observational data for structural health identification using artificial neural networks. Section 4.4 details the step-by-step construction of the AIF generative model, namely the four-tuple $\langle \mathbf{A}, \mathbf{B}, \mathbf{c}, \mathbf{d} \rangle$. Section 4.5 presents ADT simulations results under purely goal-directed behavior, which serve as a baseline for comparison with simulations involving mixed pragmatic-epistemic behavior, discussed in Section 4.6. Finally, Section 4.7 reports

Algorithm 1 Active inference loop for active digital twins.

input: generative model $\langle \mathbf{A}, \mathbf{B}, \mathbf{c}, \mathbf{d} \rangle$
 assimilated observation $O_{t_c}^{\text{Exp}} = o_{t_c}^{\text{Exp}}$
 digital state posterior $Q^*(D_{t_c-1})$ at previous time step
 action $U_{t_c-1} = u_{t_c-1}$ executed at previous time step

▷ digital state inference by minimizing variational free energy \mathcal{F}_{t_c}
 1: infer digital state posterior $Q^*(D_{t_c})$

▷ policy inference by minimizing future variational free energy $\mathcal{F}_{t_c:t_p}$
 2: compute posterior predictive distributions $Q(O_{t_c:t_p}, D_{t_c:t_p} | \pi)$
 3: evaluate epistemic and pragmatic values over $t = t_c, \dots, t_p$ under each policy
 4: infer control policies posterior $Q^*(\pi)$
 5: select action $U_{t_c} = u_{t_c}$ via best-point estimate or sampling from $Q^*(U_{t_c})$

▷ learning by minimizing variational free energy \mathcal{F}_{t_c}
 6: update transition model \mathbf{B} by updating the Dirichlet parameters $\hat{\mathbf{b}}$

return updated generative model $\langle \mathbf{A}, \mathbf{B}, \mathbf{c}, \mathbf{d} \rangle$
 updated posterior distribution over control policies $Q^*(\pi)$
 control action to be executed $U_{t_c} = u_{t_c}$
 posterior predictive density over digital states $Q(D_{t_c:t_p})$
 posterior predictive density over actions $Q(U_{t_c:t_p})$

additional results assessing the robustness of the SHM framework to incomplete data streams and to uncertainty in the extent of the damaged region.

The AIF agents based on discrete, Markovian generative models have been simulated using the open-source Python library pymdp (Heins et al., 2022). Compared to other AIF libraries, such as the MATLAB toolbox DEM (Smith et al., 2022) and the C++ library cpp-AIF (Gregoretti et al., 2024), pymdp offers notable advantages in terms of user-friendliness, flexibility, and customizability, although featuring lower process representation (DEM) and less computational efficiency (cpp-AIF). The simulations have been run on a PC featuring an Intel® Core™ i9-14900KF CPU @ 3.2 GHz and 64 GB RAM.

4.1. Physical asset

The Hörnefors railway bridge, shown in Fig. 5(a), is an integral reinforced concrete structure along the Swedish Bothnia line. It spans 15.7 m, with a clearance height of 4.7 m and a width of 5.9 m (excluding edge beams). The main structural elements have a thickness of 0.5 m for the deck, 0.7 m for the frame walls, and 0.8 m for the wing walls. The foundation system comprises two slabs connected by stay beams, supported by pile groups. The concrete is of grade C35/45, characterized by the following material properties: Young's modulus $E = 34$ GPa, Poisson's ratio $\nu = 0.2$, and density $\rho = 2500$ kg/m³. The bridge supports a single railway track with sleepers spaced at 0.65 m intervals, resting on a ballast layer that is 0.6 m deep and 4.3 m wide, with a density of $\rho_B = 1800$ kg/m³. The structure is subjected to dynamic loading from *Gröna Tåget* trains operating at speeds between $v \in [160, 215]$ km/h. We specifically consider configurations involving two-car trainsets, totaling eight axles, with each axle bearing a mass of $\psi \in [16, 22]$ ton. The geometrical and mechanical parameters, as well as the moving load model, are adapted from Rosafalco et al. (2021). The physical state space S represents the ground-truth variability in the bridge structural health.

4.2. Offline data assembly

The bridge monitoring system provides displacement data in the form of multivariate time series, denoted as $\mathbf{U}(\boldsymbol{\mu}) = [\mathbf{u}_1(\boldsymbol{\mu}), \dots, \mathbf{u}_{N_s}(\boldsymbol{\mu})] \in \mathbb{R}^{L \times N_s}$. These consist of $N_s = 10$ individual time series corresponding to the degrees of freedom (dofs) indicated in Fig. 5(b). Each series contains

L samples equally spaced over the time interval $[0, 1.5]$ s, acquired with a sampling frequency of 400 Hz. The vector $\boldsymbol{\mu} \in \mathbb{R}^{N_{\text{par}}}$ collects N_{par} system parameters, which are assumed to represent the operational and damage conditions. For the problem settings we consider, each observation spans a relatively short time interval, within which these conditions are regarded as constant.

We simulate the monitored asset using a physics-based computational model. Specifically, the structure is modeled as a linear-elastic continuum under the assumption of linearized kinematics, and the equations of elasto-dynamics describe its dynamic response to train transits. The model is spatially discretized using linear tetrahedral finite elements, and its solution is advanced in time to generate synthetic observational data, controlled by the parameter vector $\boldsymbol{\mu}$.

The full-order model (FOM) is described in detail in Torzoni et al. (2024); here, we summarize its key features. The finite element mesh consists of elements with a nominal size of 0.8 m, refined to 0.15 m along the deck, resulting in a total of 17,292 dofs. The ballast layer is accounted for by increasing the density of the deck and edge beams to represent an equivalent mass. Embankment effects are captured using distributed springs applied along the surfaces in contact with the ground, implemented via a Robin-type boundary condition with an elastic coefficient of 10^8 N/m³. Structural damping is introduced using Rayleigh damping, calibrated to yield a 5% damping ratio in the first two vibrational modes. The dynamic response is computed over the time interval $[0, 1.5]$ s, uniformly partitioned into $L = 600$ time steps, using an implicit Newmark time integration scheme (Hughes, 2000).

Damage-induced variations in the structural dynamic response are modeled as localized reductions in effective stiffness. Assuming that each observation spans a time window short enough compared to the timescale of damage progression, the structural behavior can be treated as linear within that interval. This enables a separation of timescales between the slow evolution of damage and the structural health assessment (Eftekhari Azam and Mariani, 2018). While the precise damage mechanisms are typically confirmed through on-site inspections following early detection, the degradation patterns in integral bridges that can be described in this way include: cracking in concrete due to thermal gradients, freeze-thaw cycles, or overloading; progressive deterioration from alkali-silica reactions, which may lead to cracking and spalling; cracking from stress concentrations caused by differential settlements; and surface erosion from prolonged environmental exposure.

The digital state space \mathcal{D} includes a set of predefined configurations of damage presence, location, and severity. These are modeled by parametrizing the stiffness matrix using two variables $y \in \mathbb{N}$ and $\delta \in \mathbb{R}$, both included in the parameter vector $\boldsymbol{\mu}$. The discrete variable $y \in \{0, \dots, 6\}$ designates the damage region, with $y = 0$ denoting the undamaged baseline. For the damage cases $y = 1, \dots, 6$, we consider $N_{\Omega} = 6$ predefined subdomains Ω_m , for $m = 1, \dots, 6$, each representing a potential damage location as shown in Fig. 5. Within each subdomain, the material stiffness may be reduced by a factor $\delta \in [30\%, 80\%]$, which remains constant throughout the passage of a train.

To reduce the computational cost of solving the FOM for arbitrary values of $\boldsymbol{\mu}$, we employ a projection-based reduced-order model (ROM). The reduction is performed using a Galerkin reduced basis method (Quarteroni et al., 2015; Chinesta et al., 2017), relying on a low-dimensional set of basis functions computed through proper orthogonal decomposition. Following the method of snapshots (Sirovich, 1987), the ROM is constructed upon 400 FOM solutions for different configurations of the input parameters $\boldsymbol{\mu} = (v, \psi, y, \delta)^{\top}$, which are taken as uniformly distributed and sampled via the Latin hypercube rule. The dimension of the reduced-order expansion is determined based on an energy retention criterion. By setting a tolerance of 10^{-3} for the fraction of discarded energy, the number of dofs is reduced to 133. Both the FOM and ROM have been implemented in the Matlab environment, using the redbKIT library (Negri, 2016). For a more detailed description, the reader is referred to Torzoni et al. (2024).

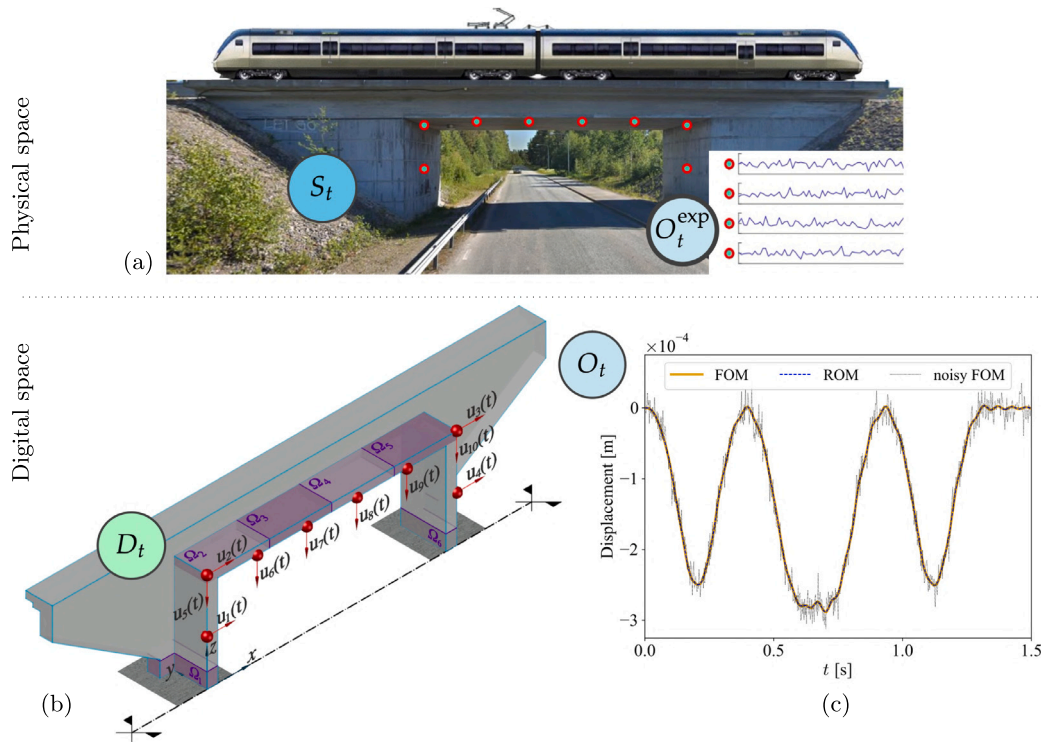


Fig. 5. Physical asset and its digital twin. (a) The physical space corresponds to the Hörnefors bridge. (b) The digital space represents a structural health monitoring schematization, including details of synthetic recordings related to displacements $u_1(t), \dots, u_{10}(t)$, and predefined damage regions $\Omega_1, \dots, \Omega_6$. (c) Exemplary vertical displacement time history at midspan, comparing full-order model (FOM), reduced-order model (ROM), and noisy FOM approximations.

A representative example of displacement time histories is reported in Fig. 5(c), showing the vertical displacement at midspan obtained from both the FOM and ROM. To emulate measurement noise and assess its potential impact on the handled structural response, signals are corrupted with additive Gaussian noise, yielding a signal-to-noise ratio of 120.

4.3. Data assimilation via artificial neural networks

The vibration recordings are assimilated for structural health diagnostics by leveraging the flexibility of deep learning (DL) models for SHM applications, as demonstrated in Torzoni et al. (2023b, 2022b) and Torzoni et al. (2023a).

Data-driven approaches to SHM follow a pattern recognition paradigm (Bishop, 2006), in which damage is assessed by comparing measurements with data previously collected under known structural conditions. This process relies on two key components: (i) feature selection and extraction, and (ii) statistical modeling to associate these features with specific damage patterns (Farrar et al., 2001). A major challenge lies in identifying damage-sensitive features that remain robust under varying operational and environmental conditions. DL offers an automated alternative for selecting and extracting optimized features by capturing temporal correlations within and across time series data (Avci et al., 2021; Fink et al., 2020).

In our framework, each time a train crosses the bridge, the vibration recordings U are initially processed by a DL classifier, which outputs confidence scores indicating the likelihood that U corresponds to each damage class defined by the y parameter. The class with the highest confidence is selected as the best-point estimate for categorizing the measurements. Whenever damage is detected and localized within a region $\Omega_m, m = 1, \dots, 6$, the vibration recordings U are further processed by a dedicated regression model — one for each damageable region — to estimate the severity of damage δ . These initial estimates are

then incorporated into the AIF framework as assimilated observation $O_{t_c}^{\text{Exp}} = o_{t_c}^{\text{Exp}}$, as detailed below.

The DL architectures have been implemented through the Tensorflow-based Keras API (Chollet et al., 2015), and trained on a single Nvidia GeForce RTX™ 3080 GPU card. The training has been performed in a supervised fashion using 10,000 noisy data instances generated from ROM simulations. For a comprehensive description the reader is referred to Torzoni et al. (2024).

4.4. Active digital twin framework

The outcomes from the DL models are integrated into our POMDP framework by discretizing the range of δ into $N_\delta = 6$ intervals: $\{[30\%, 35\%], [35\%, 45\%], [45\%, 55\%], [55\%, 65\%], [65\%, 75\%], [75\%, 80\%]\}$. This discretization results in a total of $N_\Omega N_\delta + 1 = 37$ possible damage scenarios, each specifying a combination of damage location and severity. By ordering them first by location and then by severity, this post-processed output constitutes the first observation modality $O^{\Omega\delta}$.

The observation space \mathcal{O} is completed with a second observation modality O^u corresponding to the action taken prior to data assimilation, such that $O = \{O^{\Omega\delta}, O^u\}$. Including this additional perceptual channel provides two key benefits. First, it enables prior preferences (via the c array) to account not only for the costs associated with structural health states but also for those linked to actions. Second, as detailed below, it naturally supports the formulation of an action-conditioned observation model, introducing an inductive bias that facilitates digital state identification.

The digital state space D is structured into three factors $D = \{D^{\Omega}, D^\delta, D^{\text{Epi}}\}$, corresponding to: the damage location $D^\Omega = \{\Omega_1, \dots, \Omega_6\}$; the discretized percentage reduction in material stiffness $D^\delta = \{0\%, [30\%, 35\%], [35\%, 45\%], [45\%, 55\%], [55\%, 65\%], [65\%, 75\%], [75\%, 80\%]\}$; and an epistemic switch $D^{\text{Epi}} = \{\text{Epi}, \text{Non-Epi}\}$, which

indicates whether the AIF agent is likely to engage in information-seeking (epistemic) behavior. This third factor allows the agent to autonomously switch between acting as an active information seeker or as a utility maximizer that is confident in its beliefs. It is worth noting that this factorization is neither the only viable option nor necessarily the most appropriate. This reflects the inherent subjectivity involved in shaping the digital state space. For example, D^Ω and D^δ could have been merged into a single enumerated representation, similar to the one used for $O^{\Omega\delta}$, at the expense of increased computational complexity and reduced interpretability. Alternatively, a more expressive but computationally demanding option would involve defining six separate D^δ factors, one for each of the N_Ω damageable regions. Finally, note that including D^{Epi} is essential to enable epistemic behavior, as this factor leads to distinct observation models associated with the Epi and Non-Epi states, as discussed further below.

The action space \mathcal{U} comprises four control actions, each producing specific effects:

1. *Do nothing* (DN): the structural health state evolves according to a stochastic deterioration process, while regular revenue is maintained.
2. *Maintenance* (MA): a high-cost maintenance intervention is executed to mitigate existing damage. Although this action improves the structural condition, it may not fully restore the system to a pristine (damage-free) state.
3. *Restrict operations* (RO): traffic is limited to lightweight trains with axle load below 18 ton, thereby reducing the rate of structural degradation. However, this also leads to a reduction in the revenue generated by the infrastructure.
4. *Read sensors* (RE): a moderate-cost, high-fidelity sensing action is performed to resolve uncertainty in the structural health state. This action provides high epistemic value by decreasing the entropy of the digital state posterior, thus increasing the mutual information between latent states and expected observations. This effect reflects the use of high-quality sensors, controlled forced vibration tests, or in-situ inspection. From the perspective of the generative model, performing an inspection is equivalent to reading vibration recordings from sensors, albeit with significantly higher information content and a corresponding higher cost.

The observation likelihood array $\mathbf{A} = \{\mathbf{A}^{\Omega\delta}, \mathbf{A}^u\}$ comprises two observations models: $\mathbf{A}^{\Omega\delta} \in \mathbb{R}^{|\Omega^{\Omega\delta}| \times |D^\Omega| \times |D^\delta| \times |D^{\text{Epi}}|}$ and $\mathbf{A}^u \in \mathbb{R}^{|\Omega^u| \times |D^\Omega| \times |D^\delta| \times |D^{\text{Epi}}|}$, respectively encoding the conditional sensory likelihoods $p(O^{\Omega\delta} | D^\Omega, D^\delta, D^{\text{Epi}})$ and $p(O^u | D^\Omega, D^\delta, D^{\text{Epi}})$ for the first and second observation modalities. Conceptually, these tensors are designed to answer two distinct questions: (i) what might the agent believe about the pre-classified signals? and (ii) what might the agent infer about its previous action?

The slice of $\mathbf{A}^{\Omega\delta}$ for the epistemic state, i.e., $p(O^{\Omega\delta} | D^\Omega, D^\delta, D^{\text{Epi}} = \text{Epi})$, is denoted by $\mathbf{A}_{\text{Epi}}^{\Omega\delta} \in \mathbb{R}^{|\Omega^{\Omega\delta}| \times |D^\Omega| \times |D^\delta|}$. This observation model is derived from a confusion matrix that quantifies the offline (expected) performance of the DL models in identifying the digital state factors D^Ω and D^δ . The confusion matrix is interpreted as a CPT, where rows correspond to ground-truth responses and columns to predicted outcomes. The offline evaluation has been performed using 4000 noisy FOM solutions, achieving a classification accuracy of 91.39%. To mitigate the risk of inconsistencies due to zero-likelihood observations, i.e., evidence contradicting the confusion matrix, a small positive perturbation 10^{-5} is added to all entries of $\mathbf{A}_{\text{Epi}}^{\Omega\delta}$ prior to normalization. An exemplary slice of $\mathbf{A}_{\text{Epi}}^{\Omega\delta}$ associated with $p(O^{\Omega\delta} | D^\Omega = \Omega_4, D^\delta, D^{\text{Epi}} = \text{Epi})$ is shown in Fig. 6(a).

While $\mathbf{A}_{\text{Epi}}^{\Omega\delta}$ serves as a relatively informative sensory likelihood, a higher-entropy likelihood is used to model the slice of $\mathbf{A}^{\Omega\delta}$ under the non-epistemic state, i.e., $p(O^{\Omega\delta} | D^\Omega, D^\delta, D^{\text{Epi}} = \text{Non-Epi})$, denoted by $\mathbf{A}_{\text{Non-Epi}}^{\Omega\delta} \in \mathbb{R}^{|\Omega^{\Omega\delta}| \times |D^\Omega| \times |D^\delta|}$. This non-epistemic model is obtained via uniform random perturbation of $\mathbf{A}_{\text{Epi}}^{\Omega\delta}$ as the following

linear combination:

$$\mathbf{A}_{\text{Non-Epi}}^{\Omega\delta} = (1 - \alpha)\mathbf{A}_{\text{Epi}}^{\Omega\delta} + \alpha\mathbf{A}_{\text{Entropic}}^{\Omega\delta}, \quad (24)$$

which is then properly renormalized. Here, $\mathbf{A}_{\text{Entropic}}^{\Omega\delta}$ is a purely entropic observation model sampled from a uniform distribution over $[0, 1]$, and $0 \leq \alpha \leq 1$ is a weighting coefficient controlling the degree of entropy introduced. Fig. 6(b) shows an exemplary slice corresponding to $p(O^{\Omega\delta} | D^\Omega = \Omega_4, D^\delta, D^{\text{Epi}} = \text{Non-Epi})$ for $\alpha = 0.2$. It is worth noting that modulating α can also be interpreted as a simple yet effective mechanism to account both for potential errors in, and for the decision-maker confidence about, the use of DL models to assimilate real-world data.

The slice of \mathbf{A}^u for the epistemic state, i.e., encoding $p(O^u | D^\Omega, D^\delta, D^{\text{Epi}} = \text{Epi})$, is denoted as $\mathbf{A}_{\text{Epi}}^u \in \mathbb{R}^{|\Omega^u| \times |D^\Omega| \times |D^\delta|}$. It is populated with Dirac delta distributions centered at the RE action for all possible combinations of D^Ω and D^δ (see also Fig. 6(c)). This design reflects the assumption that if the agent is in the state $D^{\text{Epi}} = \text{Epi}$, it knows with certainty that the previously taken action was the (epistemic) RE action, regardless of the values of the other digital state factors. From a data assimilation point of view, receiving $O^u = \text{RE}$ provides no informative cues for inferring D^Ω or D^δ , but it deterministically sets $D^{\text{Epi}} = \text{Epi}$. In contrast, under the non-epistemic state, the corresponding observation model $\mathbf{A}_{\text{Non-Epi}}^u \in \mathbb{R}^{|\Omega^u| \times |D^\Omega| \times |D^\delta|}$ is filled with entries that reflect a plausible causality for what the agent can infer about the previous action given D^Ω and D^δ . This prior CPT (see also Fig. 6(d)) is modeled consistently across the D^Ω factor, as follows:

$$p(O^u | D^\Omega = \Omega_{1:6}, D^\delta, D^{\text{Epi}} = \text{Non-Epi}) = \begin{pmatrix} 0.08 & 0.3 & 0.45 & 0.4 & 0.3 & 0.2 & 0.1 \\ 0.9 & 0.4 & 0.3 & 0.1 & 0.15 & 0.2 & 0.25 \\ 0.02 & 0.3 & 0.25 & 0.5 & 0.55 & 0.6 & 0.65 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}. \quad (25)$$

For data assimilation, observing $O^u \neq \text{RE}$ has two implications: first, it deterministically sets $D^{\text{Epi}} = \text{Non-Epi}$; second, it introduces an inductive bias by leveraging the structure of $\mathbf{A}_{\text{Non-Epi}}^u$ to condition inference on the previous action, similar to the influence of the transition model.

The transition array $\mathbf{B} = \{\mathbf{B}^\Omega, \mathbf{B}^\delta, \mathbf{B}^{\text{Epi}}\}$ comprises three sub-arrays $\mathbf{B}^f \in \mathbb{R}^{|\Omega^f| \times |D^f| \times |\mathcal{U}^f|}$, each encoding the transition dynamics $p(D_{t-1}^f | D_{t-1}^f, u_{t-1}^f; \mathfrak{B}^f)$ of a specific digital state factor $D^f \in \{D^\Omega, D^\delta, D^{\text{Epi}}\}$, conditioned on its previous state and the corresponding control factor $u^f \in \{u^\Omega, u^\delta, u^{\text{Epi}}\}$. Starting with initial priors over the transition probabilities defined by \mathfrak{B}^f , these are iteratively refined by assimilating evidence from the system response to actions, as described in Section 3.4. A graphical visualization of the initial transition models for each digital state and control factor is shown in Fig. 7. Note that these internal models do not replicate the ground-truth evolution, which remains unknown to the ADT. Moreover, assuming that digital state factors evolve independently, the control is factorized as $U = \{U^\Omega = \emptyset, U^\delta = \bar{U}, U^{\text{Epi}} = \bar{U}\}$. This reflects that D^Ω is an uncontrollable factor, with a control dimensionality of 1, while D^δ and D^{Epi} are both influenced by the same control variable $\bar{U} \in \mathcal{U} = \{\text{DN}, \text{MA}, \text{RO}, \text{RE}\}$. The set of feasible policies Π is constructed by combinatorially enumerating all possible sequences of actions from the action space \mathcal{U} over the prediction horizon $t = t_c, \dots, t_p$, resulting in a total of 4^{p-t_c} policies.

The initial Dirichlet parameters b^Ω over the categorical distribution \mathfrak{B}^Ω for the uncontrollable \mathbf{B}^Ω are selected to yield a 0.8 probability that damage stays in the same subdomain Ω_m , for $m = 1, \dots, 6$. The remaining 0.2 probability is evenly distributed across the other subdomains, reflecting a strong prior belief that damage is unlikely to move between different regions.

For the action-conditioned \mathbf{B}^δ , each action-specific slice encodes the probability of transitioning between discrete δ intervals. The diagonal entries represent the probability of remaining in the same damage state, while the lower-left and upper-right triangles denote the probabilities of deterioration and improvement, respectively. Under the DN action,

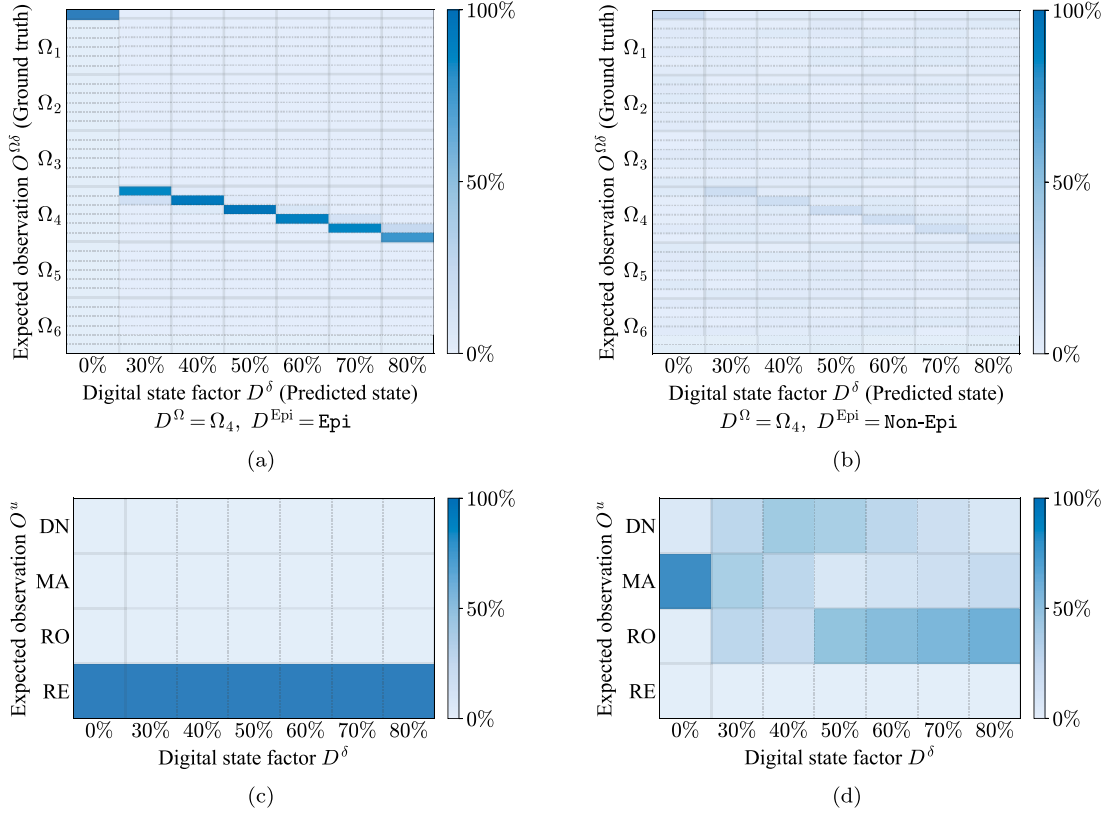


Fig. 6. Visualization of the observation models: Panels (a) and (b) show slices of A^{Ω^δ} , corresponding to the sensory likelihoods (a) $p(O^{\Omega^\delta} | D^{\Omega^\delta} = \Omega_4, D^\delta, D^{Epi} = Epi)$ and (b) $p(O^{\Omega^\delta} | D^{\Omega^\delta} = \Omega_4, D^\delta, D^{Epi} = Non-Epi)$ for $\alpha = 0.2$. Panels (c) and (d) show slices of A^u , corresponding to the sensory likelihoods (c) $p(O^u | D^{\Omega^\delta} = \Omega_{1:6}, D^\delta, D^{Epi} = Epi)$ and (d) $p(O^u | D^{\Omega^\delta} = \Omega_{1:6}, D^\delta, D^{Epi} = Non-Epi)$.

the initial Dirichlet parameters b^δ for the categorical distribution \mathfrak{B}^δ are configured to yield transition probabilities of 0.85, 0.1, and 0.05 for degradation of zero, one, or two δ intervals, respectively. For the RO action, the corresponding probabilities are set to 0.92, 0.05, and 0.03, reflecting a slower rate of deterioration due to reduced structural load. The slice associated with the RE action is designed to reflect improved damage tracking. It assigns probabilities of 0.9 and 0.1 for degradation of zero and one δ intervals, capturing the higher confidence associated with epistemic control actions. In contrast, the MA action slice is designed to support transitions across up to six δ intervals, with probabilities 0.05, 0.15, 0.20, 0.20, 0.20, and 0.20, for improvements of zero to five intervals, respectively. To mitigate the risk of numerical inconsistencies caused by evidence that contradicts the assumed transition dynamics, a small perturbation of 10^{-3} is added to all entries of B^δ prior to normalization.

The sub-array B^{Epi} serves as an epistemic switch, enabling deterministic transitions between the states $D^{Epi} = Epi$ and $D^{Epi} = Non-Epi$. This mechanism is implemented through Boolean matrices that enforce $D^{Epi} = Non-Epi$ — regardless of its previous value — whenever the ADT selects DN, MA, or RO actions. Conversely, selecting the RE action triggers a transition to the epistemic state $D^{Epi} = Epi$. This transition model is not subject to learning updates, as its structure is predefined and not expected to benefit from interaction with the generative process.

At each time step, the ADT selects a control action $u_t \in \mathcal{U}$ whose effects on the generative process are uncertain and may lead to unexpected outcomes. The costs associated with both the structural health state and the control actions are modeled as prior preferences via the array $\mathbf{c} = \{\mathbf{c}^{\Omega^\delta}, \mathbf{c}^u\}$. The components $\mathbf{c}^{\Omega^\delta} \in \mathbb{R}^{|\Omega^\delta|}$ and $\mathbf{c}^u \in \mathbb{R}^{|\mathcal{U}|}$ assign relative log-probabilities to each outcome of the two observation

modalities, respectively:

$$\mathbf{c}^{\Omega^\delta} \leftarrow \ln \tilde{p}(O_t^{\Omega^\delta}) = \begin{cases} 0 & \text{if } y = 0, \\ -\exp(\delta) & \text{if } 30\% < \delta < 80\%, \\ -10 & \text{if } \delta = 80\%, \end{cases} \quad (26)$$

$$\mathbf{c}^u \leftarrow \ln \tilde{p}(O_t^u) = \begin{cases} +5.5 & \text{if } u_t = DN, \\ -5 & \text{if } u_t = MA, \\ +2.5 & \text{if } u_t = RO, \\ -0.5 & \text{if } u_t = RE. \end{cases}$$

These log-probability vectors are passed through a Softmax function to produce valid probability distributions $\tilde{p}(O_t^{\Omega^\delta})$ and $\tilde{p}(O_t^u)$, which are then used to compute the expected utility term in the EFE. The structural health preferences penalize deterioration in proportion to the exponential of δ , with a steep penalty for severely compromised states. The control action preferences reflect trade-offs between epistemic value of expected information gain and operational cost: DN and RO actions yield positive rewards but carry the risk of structural deterioration; RE similarly allows for deterioration, yet it is expected to reduce the entropy of the digital state posterior at the cost of a moderately negative reward. MA mitigates deterioration but carries a significantly negative reward due to its high cost. While these values are expressed in non-dimensional form, they represent indicative costs charged to the decision maker. Actual values may be derived from service and cost catalogs issued by governmental agencies or infrastructure operators. In particular, the health-related preference distribution $\tilde{p}(O_t^{\Omega^\delta})$ should reflect a prioritization analysis that accounts for both the likelihood and consequences of different damage scenarios — such as loss of

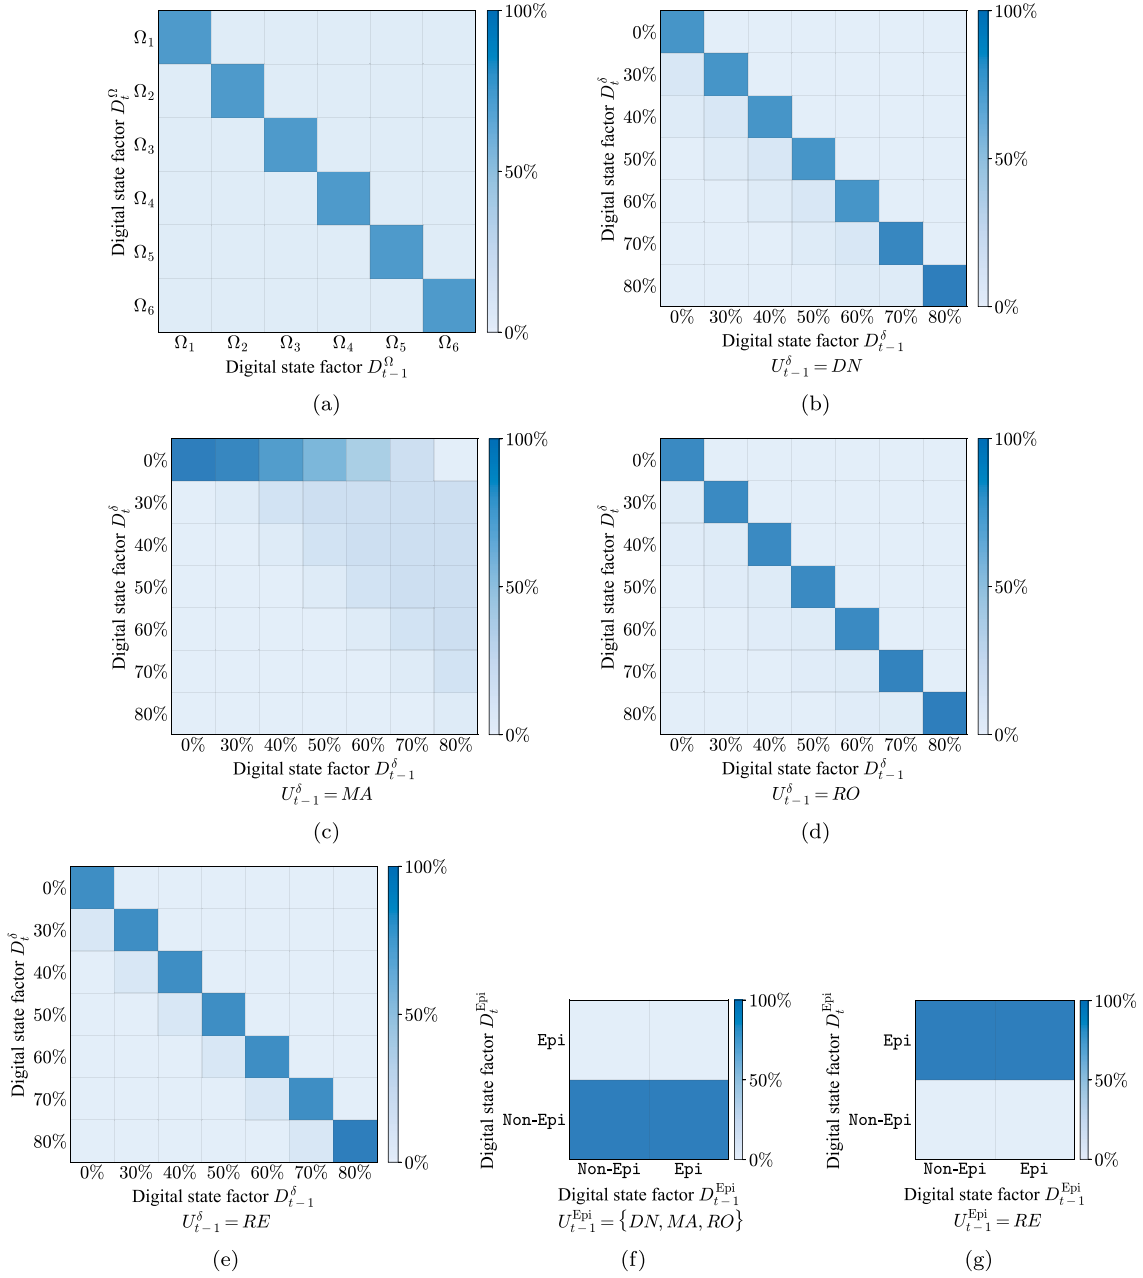


Fig. 7. Visualization of the transition models: Panel (a) shows the uncontrollable \mathbf{B}^Ω , corresponding to the transition likelihood $p(D_t^\Omega | D_{t-1}^\Omega)$. Panels (b–e) show the action-specific slices of \mathbf{B}^δ , corresponding to the transition likelihoods (b) $p(D_t^\delta | D_{t-1}^\delta, U_{t-1}^\delta = \text{DN})$, (c) $p(D_t^\delta | D_{t-1}^\delta, U_{t-1}^\delta = \text{MA})$, (d) $p(D_t^\delta | D_{t-1}^\delta, U_{t-1}^\delta = \text{RO})$, and (e) $p(D_t^\delta | D_{t-1}^\delta, U_{t-1}^\delta = \text{RE})$. Panels (f) and (g) shows the action-specific slices of \mathbf{B}^{Epi} , corresponding to the transition likelihoods (f) $p(D_t^{\text{Epi}} | D_{t-1}^{\text{Epi}}, U_{t-1}^{\text{Epi}} = \{\text{DN, MA, RO}\})$ and (g) $p(D_t^{\text{Epi}} | D_{t-1}^{\text{Epi}}, U_{t-1}^{\text{Epi}} = \text{RE})$.

serviceability, increased accident risk, or structural failure — as well as the risk tolerance of the decision-maker.

The array defining the initial state model $\mathbf{d} = \{\mathbf{d}^\Omega, \mathbf{d}^\delta, \mathbf{d}^{\text{Epi}}\}$ consists of three sub-arrays $\mathbf{d}^f \in \mathbb{R}^{|D^f|}$, each specifying the initial prior distribution $p(D_t^f)$ over a digital state factor $D^f \in \{D^\Omega, D^\delta, D^{\text{Epi}}\}$. Uniform probability distributions are adopted for D^Ω and D^{Epi} to reflect initial uncertainty. In contrast, D^δ is initialized as a Dirac delta distribution centered at 0%, consistent with the assumption of undamaged structure when the ADT enters into operation.

The (unknown) ground-truth generative process evolves conditionally on the most recent control action. In particular, we assume that damage can develop in any predefined region, without propagating across different damageable subdomains. The evolution follows the

degradation (or improvement) stochastic models described below. Under the DN, RO, and RE actions, structural health is assumed to degrade monotonically. For the DN action, the damage class y is sampled from a categorical distribution $y \sim \text{Cat}(\frac{1}{2}, \frac{1}{12}, \frac{1}{12}, \frac{1}{12}, \frac{1}{12}, \frac{1}{12}, \frac{1}{12})$, which assigns half of the probability mass to the undamaged state $y = 0$, and distributes the remaining half uniformly among the six damage classes $y = 1, \dots, 6$. When damage first initiates, the magnitude δ is sampled uniformly within the range of the first damage interval $\delta_t | y_t \neq 0, y_{t-1} = 0 \sim \text{Uniform}(0.3\%, 0.35\%)$. Subsequent damage progression is modeled by sampling δ increments from a truncated normal distribution centered at 1.5% with a standard deviation of 1%, $\delta_t - \delta_{t-1} | y_{t-1} \neq 0 \sim \text{Normal}_{\geq 0}(1.5\%, 1\%)$, with any increments below 0% rounded up to 0%. For the RO action, a similar model

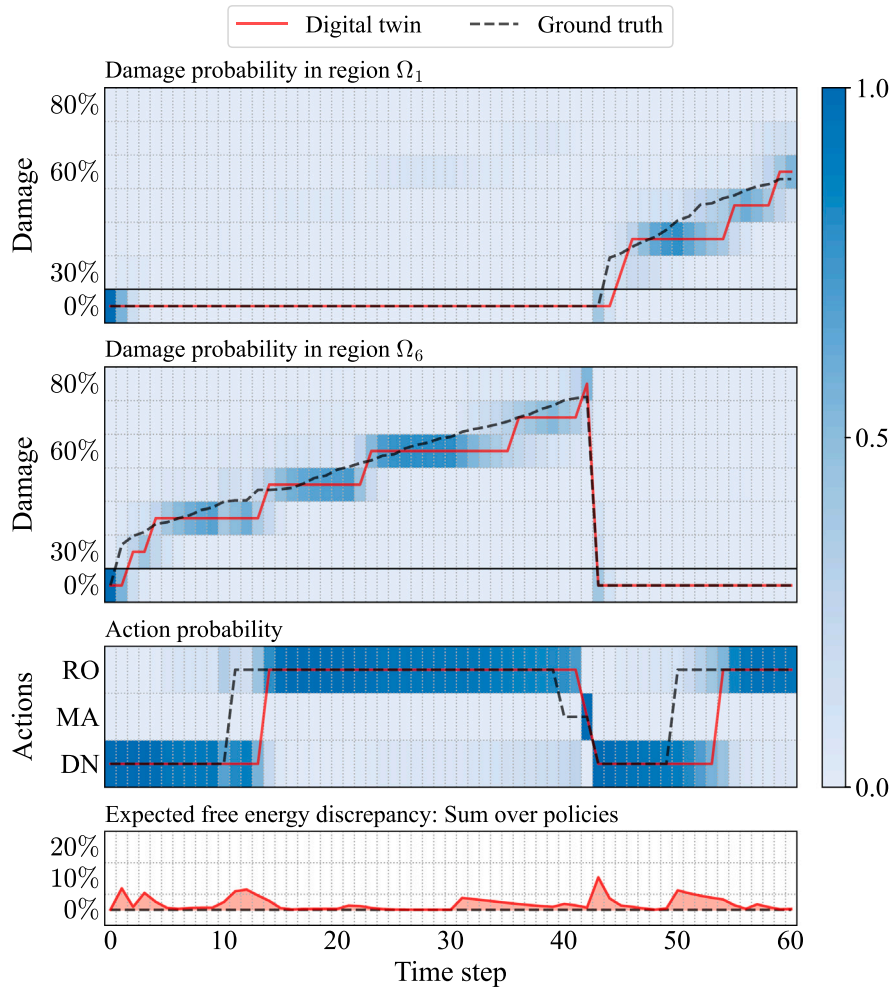


Fig. 8. Active digital twin using purely goal-directed (pragmatic) behavior. Probabilistic and best-point estimates of: (top two panels) digital state evolution compared to the ground-truth physical state; (penultimate panel) control actions recommended by the digital twin versus the optimal action under the ground-truth generative process. In the top panels, background colors represent the belief distribution over the digital state at each time step. In the penultimate panel, background colors indicate the belief distribution over the control actions. The bottom panel quantifies simulation quality in terms of the percentage absolute discrepancy between the sum of the policy-specific expected free energies computed by the digital twin and those obtained under the ground truth.

is employed, but with a lower probability of damage initiation and slower deterioration. In this case, the damage class is sampled as $y \sim \text{Cat}(\frac{3}{4}, \frac{1}{24}, \frac{1}{24}, \frac{1}{24}, \frac{1}{24}, \frac{1}{24}, \frac{1}{24}, \frac{1}{24})$, and the damage magnitude evolves as $\delta_t - \delta_{t-1} \mid y_{t-1} \neq 0 \sim \text{Normal}_{\geq 0}(0.95\%, 0.5\%)$. For the RE action, the generative process is the same as that under the DN or RO actions, respectively, depending on whether the system was previously in a restricted or unrestricted condition before engaging in information-seeking (epistemic) behavior. In contrast, the MA action is modeled as a healing process. If $y = 0$, the system remains undamaged. If $y \neq 0$, the damage magnitude decreases according to $\delta_t - \delta_{t-1} \mid y_t \neq 0 \sim \text{Normal}_{\leq 10\%}(-25\%, 15\%)$, with any decrement below 10% rounded up to 10%. The system is assumed to return to an undamaged condition ($y = 0$) if the resulting damage magnitude satisfies $\delta < 30\%$, reflecting a minimal detectable deterioration threshold.

4.5. Results: Purely goal-directed behavior

In this section and the next, we present the results of several ADT simulations, each spanning 60 time steps. At every time step, new observational data are generated based on the (unknown) ground-truth generative process. The ADT assimilates these data to infer the variational posterior $Q^*(D_{t_c})$ over the current digital state, and performs policy inference by computing the posterior $Q^*(\pi)$ over policies.

Control actions are subsequently selected as the best-point estimate from the posterior $Q^*(U_{t_c})$ over control states, and the generative model is eventually learned by updating the variational posterior Dirichlet parameters \hat{b}^* .

We adopt a policy horizon of $t_p - t_c = 4$ and begin by analyzing a baseline scenario where the ADT operates under a purely goal-directed (pragmatic) behavior. This is achieved by retaining only the utility term associated with pragmatic value in the EFE formulation, excluding any contributions from information-seeking (epistemic) value and removing the epistemic RE action from the available action set. The entropy level in the observation model $A_{\text{Non-Epi}}^{\Omega\delta}$ is set using $\alpha = 0.5$. The inverse temperature parameter controlling the precision of policy selection is left to its default value of $\gamma = 16$. Furthermore, learning updates to the generative model are disabled in this baseline setting.

Fig. 8 illustrates a representative ADT simulation. Results are reported in terms of both the ground-truth physical state and the corresponding ADT estimates obtained after assimilating observational data. The evolution of the digital state is shown only for regions that experience damage, although all damageable regions $\Omega_1, \dots, \Omega_6$ are susceptible to degradation. Initially, damage develops in Ω_6 , and the posterior $Q^*(D_{t_c})$ reveals relatively high uncertainty, primarily due to the entropy in the observation model. Nevertheless, despite the severely corrupted observation model $A_{\text{Non-Epi}}^{\Omega\delta}$, the ADT successfully follows

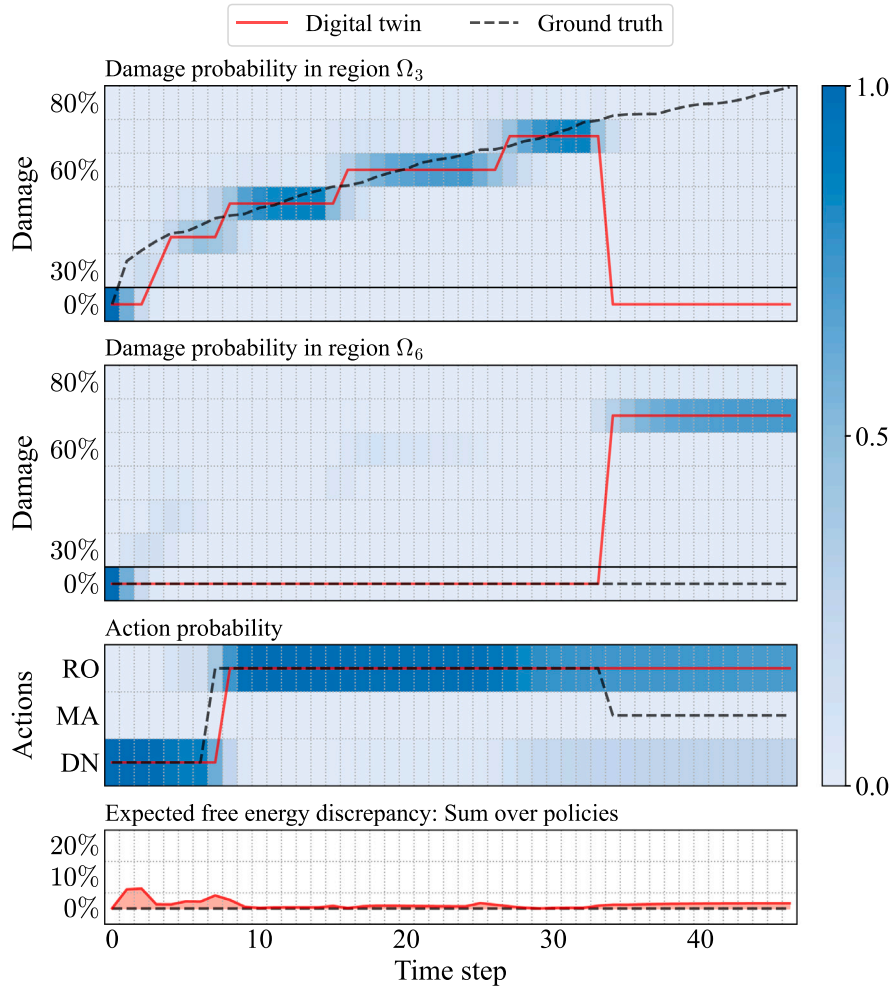


Fig. 9. Failed active digital twin using purely goal-directed (pragmatic) behavior. Probabilistic and best-point estimates of: (top two panels) digital state evolution compared to the ground-truth physical state; (penultimate panel) control actions recommended by the digital twin versus the optimal action under the ground-truth generative process. In the top panels, background colors represent the belief distribution over the digital state at each time step. In the penultimate panel, background colors indicate the belief distribution over the control actions. The bottom panel quantifies simulation quality in terms of the percentage absolute discrepancy between the sum of the policy-specific expected free energies computed by the digital twin and those obtained under the ground truth.

the ground-truth evolution by leveraging prior information from the forward-time predictor **B**. The corresponding sequence of control action estimates $Q^*(U_{t_c})$ is shown in the penultimate panel. The ADT initially recommends DN actions, aligned with the prior preferences over the two observation modalities encoded in \mathbf{c} , i.e., to maximize utility. Once a substantial probability mass in $Q^*(D_{t_c})$ is assigned to $D^\delta \geq 45\%$, RO actions begin to be selected, enabling the ADT to continue monitoring degradation, which now evolves at a reduced rate. Eventually, an MA action is selected when the structural state becomes critically compromised, as indicated by a consistent probability mass over $D^\delta \geq 75\%$ in $Q^*(D_{t_c})$. A similar behavior is shown for the subsequent damage event in Ω_1 .

For comparison, control actions under the ground-truth generative process are computed using a second AIF agent that mirrors the ADT architecture but has access to the true physical state. The ADT selects the appropriate control action from $Q^*(U_{t_c})$ with a delay of at most five time steps relative to the ground-truth-informed agent. This delay is mainly attributed to the continued use of a highly entropic observation model, which limits fast and accurate inference, along with the need to recursively update prior beliefs from earlier time steps. Note that including the RE action in the available action set would not affect the results in this case, as the ADT is driven solely by (pragmatic) utility

maximization and does not engage in information-seeking (epistemic) behavior, i.e., it does not seek to reduce the entropy of $Q(D_{t_c:t_p})$ through exploratory actions. The bottom panel of the figure assesses simulation quality by tracking the evolution of the percentage absolute discrepancy between the sum of the policy-specific EFEs computed by the ADT and those obtained under the ground-truth-informed agent:

$$\Delta^G = \left| \frac{\sum_{\pi \in \Pi} (G^\pi - \hat{G}^\pi)}{\sum_{\pi \in \Pi} \hat{G}^\pi} \right| \cdot 100, \quad (27)$$

where \hat{G}^π denotes the EFE associated with policy π under the ground-truth-informed AIF agent.

A second representative ADT simulation is shown in Fig. 9, exemplifying the same purely goal-directed (pragmatic) behavior but with a different random seed. In this case, damage begins to develop in region Ω_3 , and the ADT initially behaves consistently with the previous results, tracking the generative process with relatively high-entropy estimates propagated forward in time. However, starting from time step $t = 33$, the ADT begins to diverge from the ground truth, and the digital state posterior $Q^*(D_{t_c})$ progressively loses synchronization with the physical state. The probability mass in $Q^*(D_{t_c})$ gradually shifts from $D^\Omega = \Omega_3$ to $D^\Omega = \Omega_6$, where D^δ is consistently underestimated as lying within

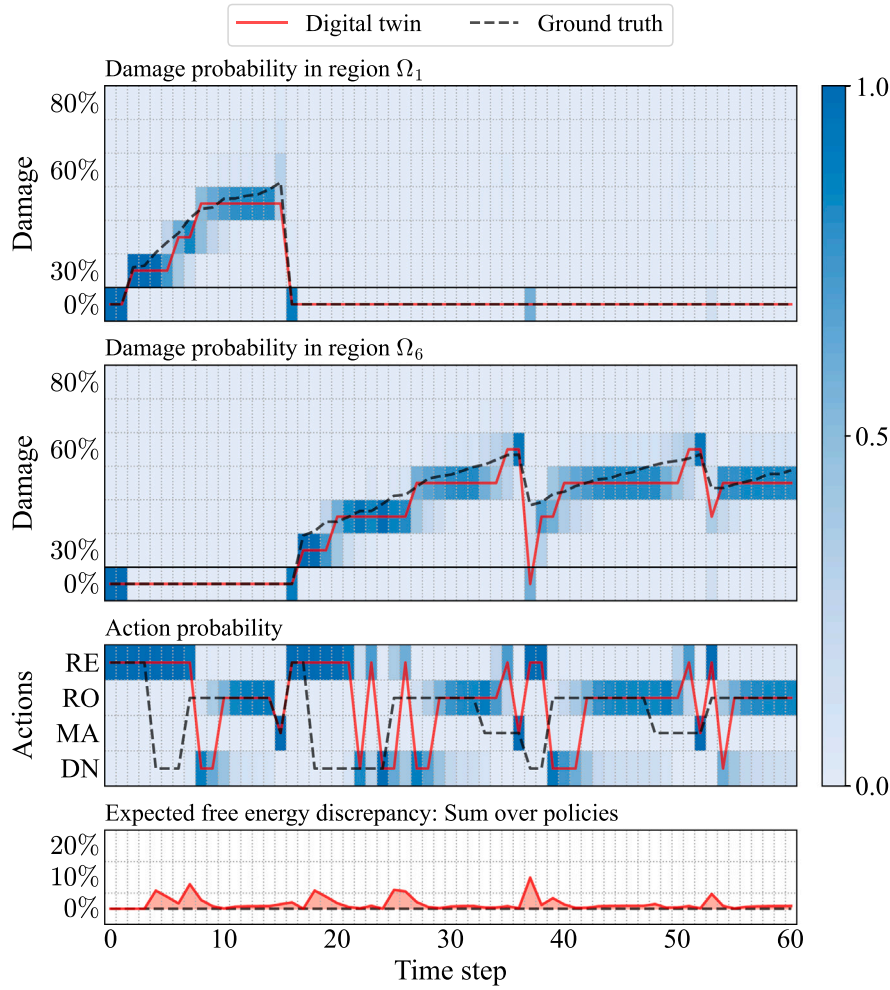


Fig. 10. Active digital twin using a combination of goal-directed (pragmatic) and information-seeking (epistemic) behaviors. Probabilistic and best-point estimates of: (top two panels) digital state evolution compared to the ground-truth physical state; (penultimate panel) control actions recommended by the digital twin versus the optimal action under the ground-truth generative process. In the top panels, background colors represent the belief distribution over the digital state at each time step. In the penultimate panel, background colors indicate the belief distribution over the control actions. The bottom panel scores simulation quality in terms of the percentage absolute discrepancy between the sum of the policy-specific expected free energies computed by the digital twin and those obtained under the ground truth.

the range [65%, 75%], while the actual value is in the range [75%, 80%]. As a result, the ADT fails to select an MA action for more than ten time steps, despite its necessity. The simulation eventually terminates at time step $t = 46$, due to a digital failure at $t = 47$, caused by $D^\delta > 80\%$, and symbolizing structural collapse. The ADT inability to recover accurate tracking is attributed to the interplay between the poorly informative sensory likelihood and the recursive propagation of outdated prior beliefs, which degrade over time.

By running a cluster of 200 simulations, each spanning 60 time steps and initialized with a different random seed for both the observation model $A_{\text{Non-Epi}}^{\Omega\delta}$ and the ground-truth generative process, the ADT operating under a purely goal-directed (pragmatic) behavior fails in 72 out of 200 cases. In this baseline setting, the failure rate is largely driven by the high entropy of the observation model $A_{\text{Non-Epi}}^{\Omega\delta}$. In-simulation performance is assessed by measuring the accuracy of the maximum a-posteriori estimate of the D^Ω and D^δ digital state factors against the ground-truth generative process. The effect of the δ discretization is accounted for by computing accuracy with a 10% tolerance (equal to the discretization step). Under these conditions, the ADT achieves a mean accuracy of 69% with a 95% confidence interval of $\pm 2.6\%$. Although highly corrupted observations reflect challenging real-world

conditions, the results presented in the following section show that equipping the ADT with both goal-directed and information-seeking (epistemic) components enables active exploration in response to critical uncertainty, resulting in a significant performance improvement over this purely pragmatic baseline.

The behavior described above can also be illustrated using a simplified scenario in which the ADT relies on two sensors, each providing partial observations to update its beliefs about the evolving damage state. If one sensor becomes faulty but the transition model closely approximates the actual dynamics of damage progression, the ADT may still track the system accurately, as the predictive power of the prior compensates for the degraded sensory evidence. However, in the more typical case where the transition model does not fully capture the actual system evolution, outdated priors dominate the inference process, and the compromised likelihood is unable to correct them. In such conditions, an information-seeking (epistemic) action should ideally be triggered to resolve ambiguity and restore confidence in the likelihood model — for instance, by querying a redundant sensor, activating a dormant one, or scheduling a targeted diagnostic procedure.

It is interesting to note how the EFE discrepancy shown in the bottom panel of Fig. 9 does not indicate any critical issue. This is because

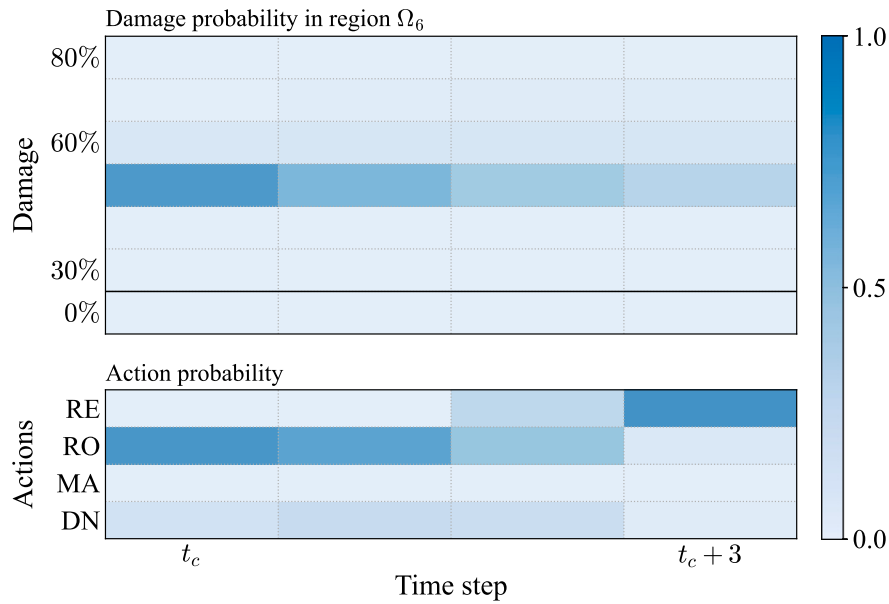


Fig. 11. Active digital twin using a combination of goal-directed (pragmatic) and information-seeking (epistemic) behaviors. Posterior predictive densities beyond data assimilation over (future) digital states and control states, starting at $t_c = 60$. In the top panel, background colors represent the belief distribution over the digital state at each time step. In the bottom panel, background colors indicate the belief distribution over the control actions.

the EFE is a subjective metric, reflecting how the ADT evaluates its own performance rather than measuring the correctness of the digital state estimates. Indeed, the Δ^G indicator quantifies the misalignment in belief-driven action planning between the ADT and an idealized agent with access to the true physical state. As a result, Δ^G remains low simply because the ADT is (mistakenly) confident in its estimated behavior, just as the ground-truth-informed agent is confident in its own. However, the resulting control actions differ. To address this limitation, one could instead employ objective performance indicators derived from the generative process. Examples include utility scores evaluated on realized system outcomes, or the survival time before reaching a critical condition.

4.6. Results: Combining goal-directed and information-seeking behaviors

In this section, we present the results of ADT simulations combining goal-directed (pragmatic) and information-seeking (epistemic) behaviors. For this, we adopt the complete EFE formulation including both pragmatic and epistemic terms; furthermore, we include the epistemic RE action in the available actions. All other settings remain unchanged from the simulations presented earlier.

Fig. 10 illustrates a representative simulation. The ADT initially exhibits information-seeking (epistemic) behavior, executing a sequence of RE actions to gather information about damage onset. Once the posterior $Q^*(D_{t_c})$ identifies evolving damage within Ω_1 with relatively low uncertainty, the ADT shifts to DN actions aimed at (pragmatic) utility maximization. When a significant portion of $Q^*(D_{t_c})$ supports $D^\delta \geq 45\%$, RO actions begin to emerge. An MA action is eventually selected when the risk of structural failure becomes substantial, i.e., for a significant probability mass over $D^\delta \geq 65\%$. A similar pattern is observed during the subsequent damage event in Ω_6 . In this case, sporadic RE actions are also triggered whenever the entropy of $Q^*(D_{t_c})$ increases, to prevent desynchronization from the physical state. These RE actions are interleaved with extended sequences of DN and RO decisions, depending on the evolving health state and the interaction between the sensory likelihood and the transition model. For instance, RE actions at $t = 35$ and $t = 51$ are deployed to disambiguate the digital state just before executing costly MA interventions. In contrast, the first MA action at $t = 15$ is not preceded by RE behavior, as the

ADT maintains a confident, low-entropy belief at that point. Note that the epistemic RE action carries a lower prior preference $\tilde{p}(O^u)$ than DN or RO actions and does not directly affect damage progression. As a result, it is employed only under epistemic-driven behavior, where its role is to support future goal-directed (pragmatic) decisions. Similarly, RE actions occurring immediately after MA interventions reflect the ADT effort to resolve uncertainty via active exploration of maintenance outcomes. In contrast, under the ground-truth generative process, RE actions are triggered exclusively to gather initial evidence about damage onset. Corrective inference is unnecessary in this setting due to perfect, uncertainty-free access to the physical state.

Fig. 11 shows the posterior predictive densities for the future digital states $Q^*(D_{t_c:t_p})$ and the corresponding control states $Q^*(U_{t_c:t_p})$, starting at $t_c = 60$ and spanning four time steps. These predictions capture the expected progression of structural health, conditioned on the posterior over policies $Q^*(\pi)$, thereby supporting the planning of preventive interventions. When belief propagation leads to an overly flat digital state distribution, the likelihood of selecting an RE action increases, mitigating the risk of decisions based on unreliable or uncertain belief states.

By running a second cluster of 200 simulations, each initialized with a different random seed, the ADT operating under combined goal-directed (pragmatic) and information-seeking (epistemic) behaviors exhibits zero failures. In-simulation performance also increases, with the maximum a-posteriori estimate of the D^Ω and D^δ digital state factors achieving a mean accuracy of 89% with a 95% confidence interval of $\pm 0.5\%$. This result underscores the potential of fully equipped ADTs compared to the purely pragmatic baseline discussed in Section 4.5.

The results of a complete ADT simulation, combining goal-directed (pragmatic) and information-seeking (epistemic) behaviors, and additionally incorporating learning updates to the generative model, are shown in Fig. 12 for the same initialization seed as in Fig. 10. This scenario spans 80 time steps and introduces learning via updates to the transition model array \mathbf{B} , with a learning rate of $\eta = 0.1$. Learning demonstrates beneficial in several aspects. First, it reduces the frequency of incorrect digital state inferences, thereby shortening the average response delay relative to the ground-truth agent. Second, as the transition dynamics become progressively tailored to the (unknown) generative process, the ADT gains confidence in its predictions,

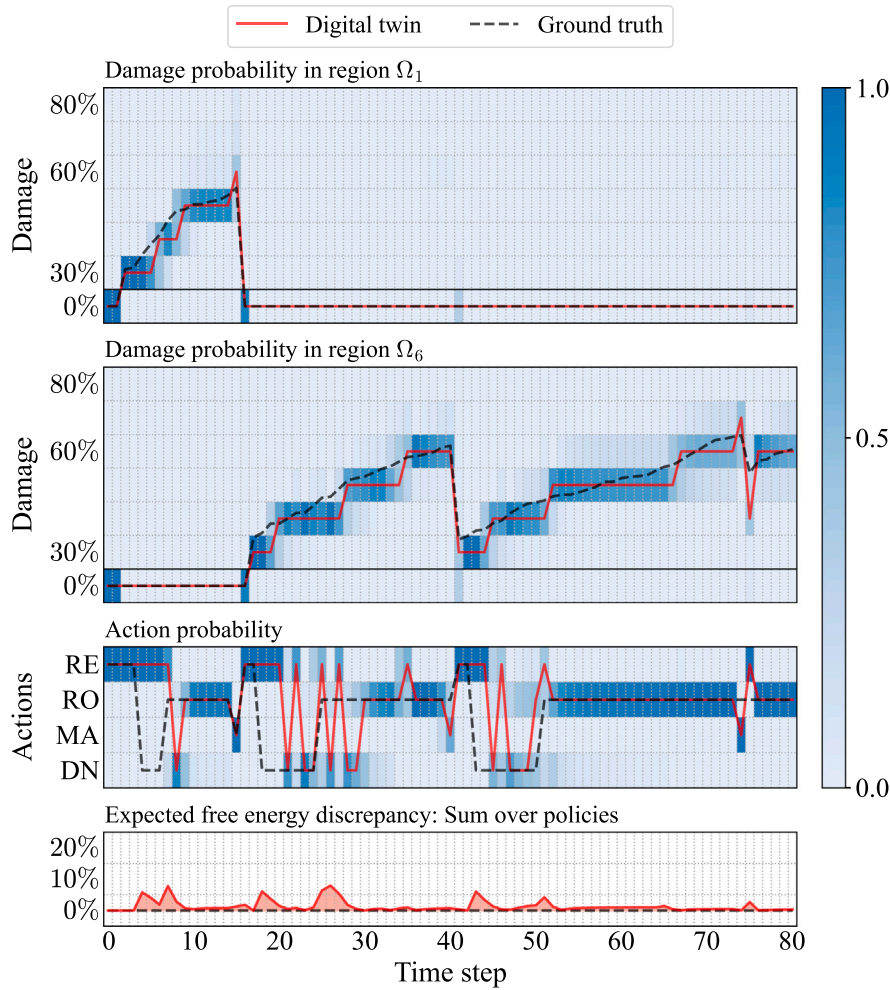


Fig. 12. Active digital twin using a combination of goal-directed (pragmatic) and information-seeking (epistemic) behaviors, and additionally incorporating learning updates to the generative model. Probabilistic and best-point estimates of: (top two panels) digital state evolution compared to the ground-truth physical state; (penultimate panel) control actions recommended by the digital twin versus the optimal action under the ground-truth generative process. In the top panels, background colors represent the belief distribution over the digital state at each time step. In the penultimate panel, background colors indicate the belief distribution over the control actions. The bottom panel quantifies simulation quality in terms of the percentage absolute discrepancy between the sum of the policy-specific expected free energies computed by the digital twin and those obtained under the ground truth.

resulting in a reduced need for (corrective) RE actions. Third, the gradual reduction of uncertainty in the transition model enables the ADT to safely delay maintenance toward the end of the simulation. For example, the third maintenance action, previously triggered at $t = 52$, is now postponed to $t = 74$. Moreover, this intervention is no longer based on a maximum a-posteriori estimate of D^δ within the [55%, 65%] range, but instead within the higher [65%, 75%] range — highlighting the potential for resource savings across the system operational lifespan. The runtime for this simulation is about 130 s, averaging 1.6 s per time slice.

4.7. Results: Robustness assessment

In this section, we assess the robustness of the ADT, fully equipped with goal-directed (pragmatic) and information-seeking (epistemic) behaviors, against incomplete data streams and uncertainty in the extent of the damaged region. We first consider two stress-testing scenarios: (i) progressively more incomplete data streams during the passage of a train, and (ii) progressively more frequent observation failures.

Incomplete data streams during the passage of a train are modeled by zeroing vibration recordings in U over the central portion of the monitoring window [0.5 s, 1 s]. Fig. 13 shows representative cases

where the corruption affects an increasing number of sensors, up to all sensors. As information loss increases, the ADT becomes more prone to desynchronize from the ground truth, with the risk of failure due to the delayed selection of a necessary MA action. To counteract this, the ADT autonomously increases the probability of selecting corrective RE actions, promoting physical–digital realignment. Interestingly, corrupting only a subset of sensors proves more detrimental than disabling all sensors over the same time window. This effect is due to the disruption of cross-channel convolution patterns in the DL models used to construct the observation modality O^{Ω^δ} . Partial corruption generates inconsistent patterns not encountered during training, whereas removing all channels leads to information loss without injecting misleading correlations. In this latter case, the ADT still exploits the remaining signal segments, although with reduced tracking accuracy and an increased need for RE actions.

From a slightly different perspective, observation failures are instead modeled by making observational evidence for O^{Ω^δ} unavailable to the ADT at selected time steps. Fig. 14 presents two representative simulations in which measurements are received only every two or five simulation steps. In both cases, the ADT remains synchronized with the ground truth by leveraging the predictive capabilities of the transition model. The main effect is a periodic widening and sharpening of the

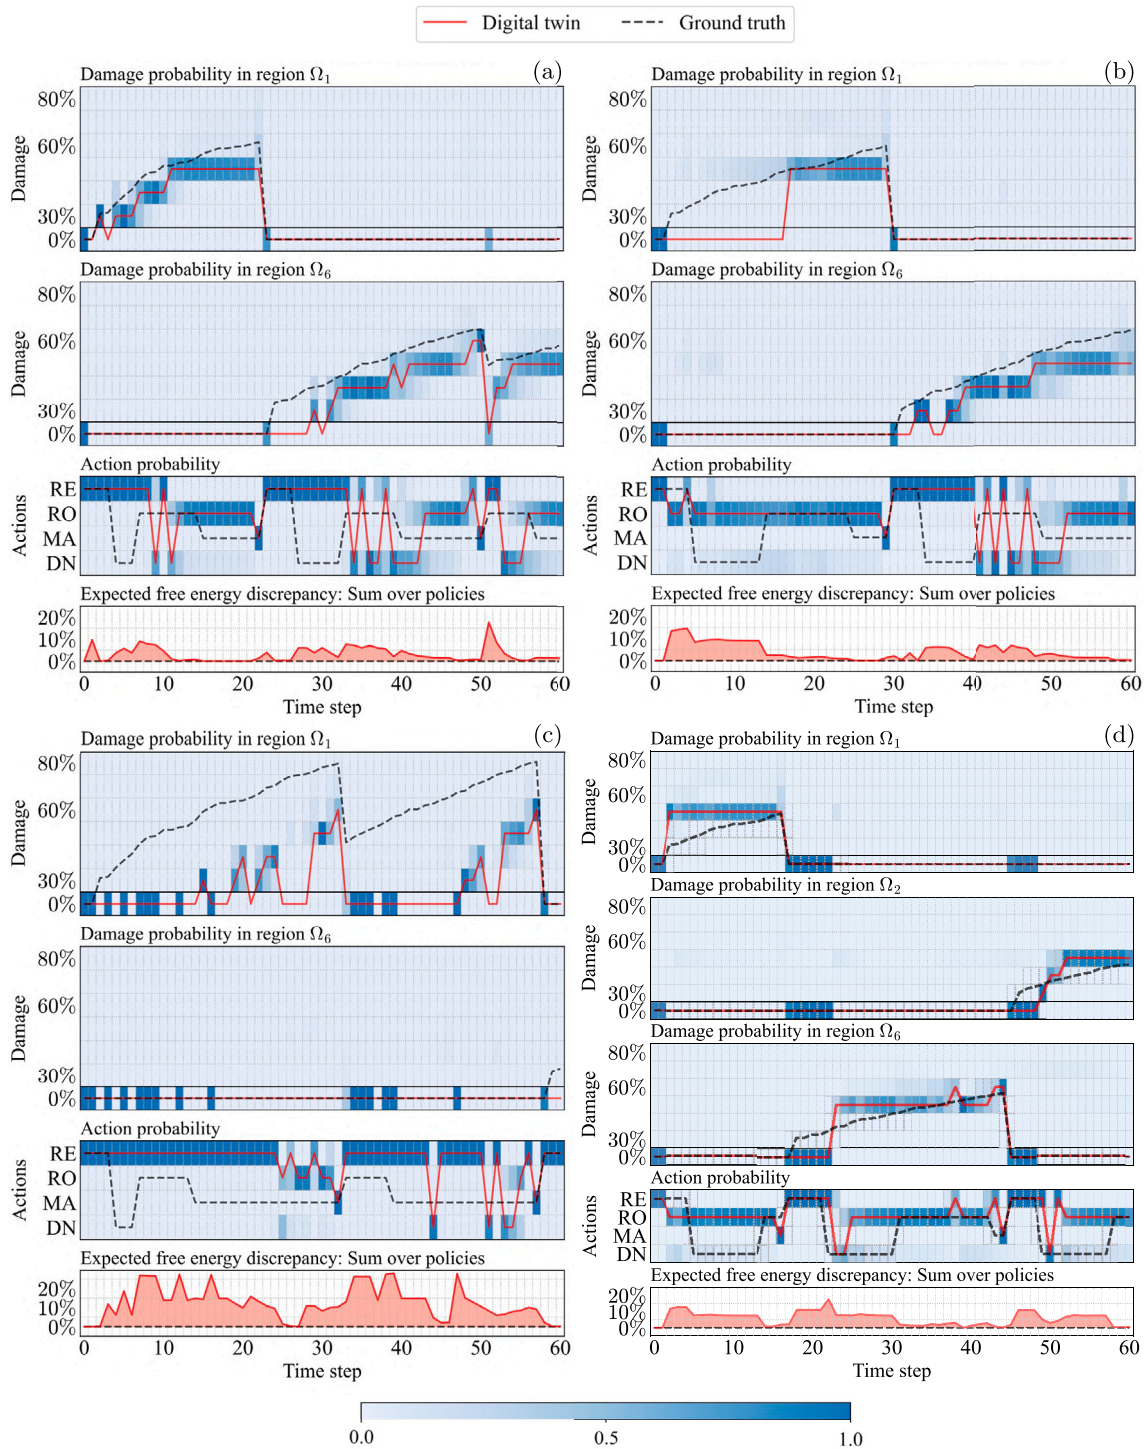


Fig. 13. Active digital twin using a combination of goal-directed (pragmatic) and information-seeking (epistemic) behaviors. Performance under progressively more incomplete data streams during the passage of a train, affecting: (a) sensors $\{u_5, u_{10}\}$; (b) sensors $\{u_3, u_5, u_{10}\}$; (c) sensors $\{u_1, u_3, u_5, u_{10}\}$; and (d) all sensors $\{u_1, \dots, u_{10}\}$. Incomplete data streams correspond to the loss of one third of the normally sampled measurements, modeled by zeroing the signals over the central portion of the monitoring window [0.5 s, 1 s]. Probabilistic and best-point estimates of: (top panels) digital state evolution compared to the ground-truth physical state; (penultimate panels) control actions recommended by the digital twin versus the optimal action under the ground-truth generative process. In the top panels, background colors represent the belief distribution over the digital state at each time step. In the penultimate panels, background colors indicate the belief distribution over the control actions. The bottom panels scores simulation quality in terms of the percentage absolute discrepancy between the sum of the policy-specific expected free energies computed by the digital twin and those obtained under the ground truth.

belief over the digital state, reflecting the observation frequency. This leads to a stronger tendency to select information-gathering RE actions rather than the more conservative RO action, despite the latter having a higher prior preference $\tilde{p}(O^u)$.

Finally, robustness to modeling inaccuracies is assessed by introducing uncertainty in the size of the damaged region. As in previous cases, test instances are derived from FOM solutions corrupted with additive Gaussian noise. However, here the extent of the damaged subdomains

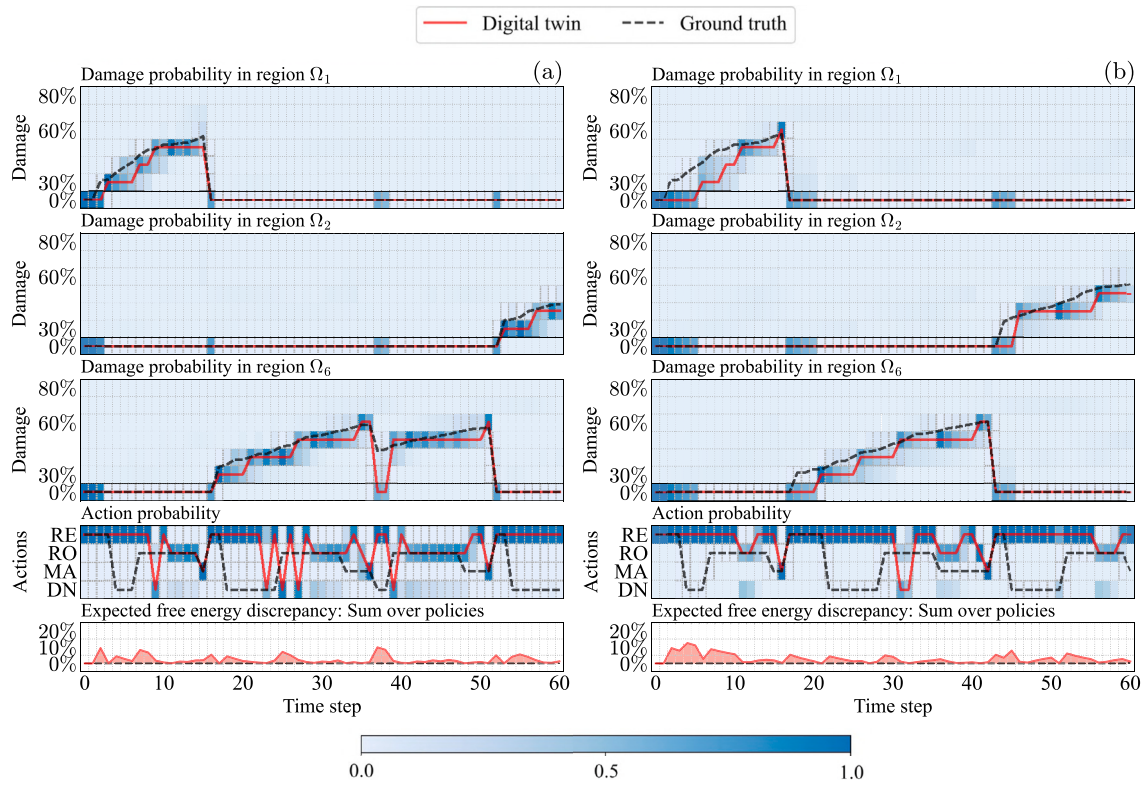


Fig. 14. Active digital twin using a combination of goal-directed (pragmatic) and information-seeking (epistemic) behaviors. Performance under progressively more frequent observation failures, with the active digital twin receiving observational data: (a) every two simulation steps; and (b) every five simulation steps. When a measurement fails, observational evidence for the first modality O^{δ} is unavailable to the active digital twin. Probabilistic and best-point estimates of: (top three panels) digital state evolution compared to the ground-truth physical state; (penultimate panels) control actions recommended by the digital twin versus the optimal action under the ground-truth generative process. In the top panels, background colors represent the belief distribution over the digital state at each time step. In the penultimate panels, background colors indicate the belief distribution over the control actions. The bottom panels scores simulation quality in terms of the percentage absolute discrepancy between the sum of the policy-specific expected free energies computed by the digital twin and those obtained under the ground truth.

$\Omega_1, \dots, \Omega_6$ is reduced by 25% and 50% in the testing data only; the offline training dataset remains unchanged.

Fig. 15 shows a representative case for a 50% reduction in the extent of Ω_6 . The ADT correctly identifies the damaged region but systematically underestimates its magnitude. The resulting behavior is consistent with what seen in Section 4.6: an initial sequence of epistemic RE actions is followed by pragmatic DN actions, interleaved with RE actions to compensate for increased $Q^*(D_{t_c})$ entropy. Subsequently, RO actions emerge to slow degradation, and a MA action is selected when failure risk becomes substantial, anticipated by a preventive RE action.

While the persistent underestimation may degrade tracking performance, it does not necessarily imply desynchronization. We quantify this risk through a cluster of 200 simulations with different random seeds for a 50% size reduction. The mean accuracy of the maximum a-posteriori estimates of D^Ω and D^δ decreases to 40% (95% confidence interval of $\pm 3.2\%$), yet only a single failure occurs. For a 25% size reduction, the mean accuracy increases to 80.5% (95% confidence interval $\pm 0.8\%$), with no observed failures. Despite the degradation in tracking accuracy, these results demonstrate robustness to misspecification of damage extent, in addition to previously considered uncertainties such as measurement noise, varying operational conditions, and damage severity.

5. Discussion

We have shown that framing a digital twin as an AIF agent transforms it from a passive observer into an autonomous decision-making

entity. The considered case study, among the few AIF applications in engineering, demonstrates how pragmatic and epistemic behaviors emerge naturally from expected free energy minimization, without explicit programming. Beyond optimizing structural health and maintenance costs, the ADT maintains synchronization with the evolving structural condition by autonomously deciding when to acquire additional information. Actions are selected not only to respond to the current belief state, but also to shape future observations and reduce anticipated uncertainty. The ADT also proves robust to incomplete or unreliable observations and to uncertainty in the extent of the damaged region, adapting its policy to mitigate desynchronization risks.

These capabilities stem directly from the advantages of AIF. Utility maximization is not the sole driver of policy selection; rather, it complements information-gain maximization within a unified functional objective that balances exploration and exploitation. Unlike reinforcement learning, which typically does not consider information-gain maximization, and relies on reward functions shaped through extensive trial-and-error over large datasets, AIF offers greater modeling and learning flexibility by encoding preferences and beliefs as probability distributions within a generative model. This feature is particularly valuable in complex, nonstationary, and nonlinear environments, where robustness and adaptivity are critical. Moreover, explicitly incorporating hypotheses and assumptions into the generative model enhances interpretability, boosting explainable and trustworthy decision-making (Albarracín et al., 2024). However, we note that reinforcement learning can be viewed as a special case of AIF (Friston et al., 2015). Since rewards or utility functions can be expressed as log-probabilities encoding prior preferences, reinforcement learning

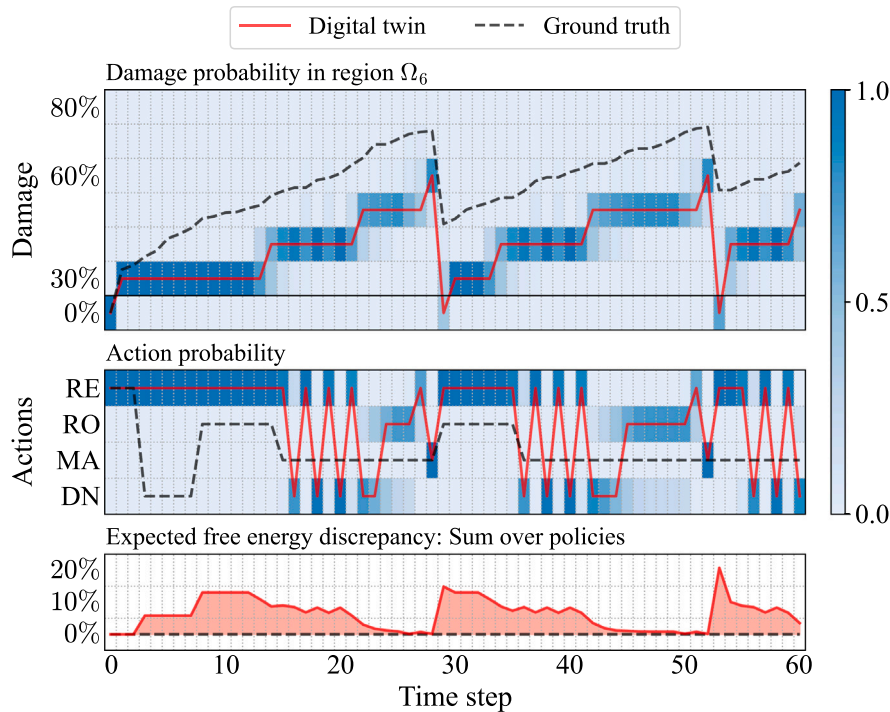


Fig. 15. Active digital twin using a combination of goal-directed (pragmatic) and information-seeking (epistemic) behaviors. Performance for a damage region of size reduced by 50% compared to the training regime. Probabilistic and best-point estimates of: (top panel) digital state evolution compared to the ground-truth physical state; (middle panel) control actions recommended by the digital twin versus the optimal action under the ground-truth generative process. In the top panel, background colors represent the belief distribution over the digital state at each time step. In the middle panel, background colors indicate the belief distribution over the control actions. The bottom panel scores simulation quality in terms of the percentage absolute discrepancy between the sum of the policy-specific expected free energies computed by the digital twin and those obtained under the ground truth.

can be reformulated as inference over policies that generate preferred outcomes.

Limitations: The effectiveness of the ADT framework depends on the design of the generative model, including the choice of the digital state space and the specification of the observation and transition models, which remain problem-dependent. In the presented application, one observation modality relies on deep learning models trained offline on synthetic data. Although this supports generalization to realistic scenarios, performance is sensitive to modeling assumptions and to the representativeness of the training dataset. Capturing more complex damage patterns or geometries would require more detailed numerical models and larger datasets, increasing computational demands. These aspects are also linked to sensor deployment and the associated value of information (Andriotis et al., 2021; Malings and Pozzi, 2018), highlighting the importance of optimal sensor placement (Capellari et al., 2018).

Policy recommendations: From an operational perspective, digital twins should incorporate decision-making and information-seeking capabilities, rather than being limited to state estimation and prediction. Moreover, probabilistic representations such as PGMs are essential to explicitly manage uncertainty and enable adaptive behavior. While ADTs enjoy these features, their behavior is not fixed but depends on the objectives encoded in the generative model. Prior preferences can be tuned to reflect specific safety or operational requirements, allowing stakeholders to explore how varying levels of risk sensitivity influence decisions. Active digital twin-based systems can thus support the transition from periodic or reactive maintenance to predictive strategies, where inspection and monitoring are scheduled autonomously based on quantified uncertainty and expected utility.

6. Conclusions and outlook

Summary of contributions: This paper introduces active digital twins based on the active inference paradigm. At the core of active digital twins lies a self-updating generative model that interacts with a partially observable dynamical environment, extending the abstraction of physical-digital systems proposed by Kapteyn et al. (2021). By leveraging the variational free energy minimization process that drives active inference agents (Parr et al., 2022; Friston, 2010), we have enabled active digital twins capable of adaptively monitoring, interacting with, and learning from uncertain and dynamic environments. Of particular interest is their ability to autonomously balance goal-directed (pragmatic) and information-seeking (epistemic) behaviors. Within this dual objective, decision-support from the active digital twin becomes an integral part of the inference process, allowing uncertainty about hidden states to be resolved as a means to maximize utility, in the spirit of intelligent automation (San et al., 2026).

Results have been presented for a case study on structural health monitoring and predictive maintenance of a railway bridge. The application has focused on the construction of a generative model enabling bidirectional perception-action interaction. Simulations of active digital twins have been carried out using active inference agents with progressively richer behaviors: purely goal-directed; combined goal-directed and information-seeking; with or without learning updates to the generative model. The results demonstrate that active digital twins autonomously balance the joint optimization of structural health and maintenance costs objectives with the need to acquire information in a principled manner, as dictated by free energy minimization. In particular, active exploration has proved essential for maintaining synchronization between the digital and physical states,

while incorporating learning updates has improved inference accuracy, reduced the need for corrective epistemic actions, and enabled the safe postponement of costly interventions.

Opportunities for future research: Beyond the structural health monitoring application presented here, the proposed framework offers a generalizable methodology applicable across a wide range of domains. Active digital twins are envisioned as key enablers of autonomous agents in the development of smart structures and systems (Zhang et al., 2023), as well as in fields such as medicine and neuroscience (Amunts et al., 2024). Future extensions can leverage the epistemic value associated with expected information gain not only over digital states but also over model parameters, as anticipated in Section 3.5, enabling the generative model underlying active digital twins to be learned online from controlled experience. This capability will support complex behaviors combining goal-directed, information-gathering, and curiosity-driven components (Friston et al., 2017a), fostering adaptation and continual self-learning from real-world data. Online learning may also be complemented by offline updates via Bayesian model reduction (Friston et al., 2018), enabling an optimal trade-off between the complexity and accuracy of the generative model.

CRedit authorship contribution statement

Matteo Torzoni: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Domenico Maisto:** Writing – review & editing, Methodology, Formal analysis, Conceptualization. **Andrea Manzoni:** Writing – review & editing, Methodology, Funding acquisition, Formal analysis, Conceptualization. **Francesco Donnarumma:** Writing – review & editing, Methodology, Formal analysis, Conceptualization. **Giovanni Pezzulo:** Writing – review & editing, Supervision, Methodology, Funding acquisition, Formal analysis, Conceptualization. **Alberto Corigliano:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Funding acquisition, Formal analysis, Conceptualization.

Funding

This work is supported by the ERC advanced grant IMMENSE (Grant Agreement 101140720), funded by the European Union. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. Author AM also acknowledges the financial support from the FIS starting grant DREAM (Grant Agreement FIS00003154), funded by the Italian Science Fund (FIS) - Ministero dell'Università e della Ricerca. Authors DM, FD, and GP acknowledge financial support from the ERC consolidator grant ThinkAhead (Grant Agreement 820213), funded by the European Union.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Matteo Torzoni reports financial support was provided by European Research Council. Andrea Manzoni reports financial support was provided by European Research Council. Alberto Corigliano reports financial support was provided by European Research Council. Domenico Maisto reports financial support was provided by European Research Council. Francesco Donnarumma reports financial support was provided by European Research Council. Giovanni Pezzulo reports financial support was provided by Government of Italy Ministry of University and Research.

Acknowledgments

The authors thank Dr. Luca Rosafalco (Politecnico di Milano) and Eng. Giacomo Mondello (Politecnico di Milano) for the insights and contributions during our discussions.

Appendix. Expected free energy derivations

In this Appendix, we provide the complete derivation of the expected free energy expressions, as adapted to the active digital twin framework from Heins et al. (2022):

Expected free energy.

$$\begin{aligned}
 G_t^\pi &= \mathbb{E}_{Q(O_t, D_t | \pi)} [\ln Q(D_t | \pi) - \ln \tilde{p}(O_t, D_t | \pi)] \\
 &= \mathbb{E}_{Q(O_t, D_t | \pi)} [\ln Q(D_t | \pi) - \ln \tilde{p}(O_t, D_t | \pi) \\
 &\quad + \underbrace{\ln Q(D_t | O_t, \pi) - \ln Q(D_t | O_t, \pi)}_{=0}] \\
 &= \mathbb{E}_{Q(O_t, D_t | \pi)} [\ln Q(D_t | \pi) - \ln Q(D_t | O_t, \pi) - \ln \tilde{p}(O_t) \\
 &\quad - \ln p(D_t | O_t, \pi) + \ln Q(D_t | O_t, \pi)] \\
 &= -\mathbb{E}_{Q(O_t, D_t | \pi)} [\ln Q(D_t | O_t, \pi) - \ln Q(D_t | \pi)] - \mathbb{E}_{Q(O_t, D_t | \pi)} [\ln \tilde{p}(O_t)] \\
 &\quad + \mathbb{E}_{Q(O_t, D_t | \pi)} [\ln Q(D_t | O_t, \pi) - \ln p(D_t | O_t, \pi)] \\
 &= -\underbrace{\mathbb{E}_{Q(O_t | \pi)} [\text{D}_{\text{KL}}[Q(D_t | O_t, \pi) \parallel Q(D_t | \pi)]]}_{\text{Epistemic value (information gain)}} - \underbrace{\mathbb{E}_{Q(O_t | \pi)} [\ln \tilde{p}(O_t)]}_{\text{Pragmatic value (utility)}} \\
 &\quad + \underbrace{\mathbb{E}_{Q(O_t | \pi)} [\text{D}_{\text{KL}}[Q(D_t | O_t, \pi) \parallel p(D_t | O_t, \pi)]]}_{\text{Expected variational approximation error } (\geq 0)}.
 \end{aligned} \tag{A.1}$$

Expected free energy with model parameters.

$$\begin{aligned}
 G_t^\pi &= \mathbb{E}_{Q(O_t, D_t, \phi | \pi)} [\ln Q(D_t, \phi | \pi) - \ln \tilde{p}(O_t, D_t, \phi | \pi)] \\
 &= \mathbb{E}_{Q(O_t, D_t, \phi | \pi)} [\ln Q(D_t, \phi | \pi) - \ln \tilde{p}(O_t, D_t, \phi | \pi) \\
 &\quad + \underbrace{\ln Q(D_t, \phi | O_t, \pi) - \ln Q(D_t, \phi | O_t, \pi)}_{=0}] \\
 &= \mathbb{E}_{Q(O_t, D_t, \phi | \pi)} [\ln Q(D_t | \pi) + \ln Q(\phi | \pi) - \ln \tilde{p}(O_t) - \ln p(D_t, \phi | O_t, \pi) \\
 &\quad + \ln Q(D_t, \phi | O_t, \pi) - \ln Q(D_t, \phi | O_t, \pi)] \\
 &= \mathbb{E}_{Q(O_t, D_t, \phi | \pi)} [\ln Q(D_t | \pi) - \ln Q(D_t | O_t, \pi) + \ln Q(\phi | \pi) \\
 &\quad - \ln Q(\phi | O_t, \pi) \\
 &\quad - \ln \tilde{p}(O_t) - \ln p(D_t, \phi | O_t, \pi) + \ln Q(D_t, \phi | O_t, \pi)] \\
 &= -\mathbb{E}_{Q(O_t, D_t, \phi | \pi)} [\ln Q(D_t | O_t, \pi) - \ln Q(D_t | \pi)] \\
 &\quad - \mathbb{E}_{Q(O_t, D_t, \phi | \pi)} [\ln Q(\phi | O_t, \pi) - \ln Q(\phi | \pi)] - \mathbb{E}_{Q(O_t, D_t, \phi | \pi)} [\ln \tilde{p}(O_t)] \\
 &\quad + \mathbb{E}_{Q(O_t, D_t, \phi | \pi)} [\ln Q(D_t, \phi | O_t, \pi) - \ln p(D_t, \phi | O_t, \pi)] \\
 &= -\underbrace{\mathbb{E}_{Q(O_t | \pi)} [\text{D}_{\text{KL}}[Q(D_t | O_t, \pi) \parallel Q(D_t | \pi)]]}_{\text{Epistemic value (digital state information gain)}} \\
 &\quad - \underbrace{\mathbb{E}_{Q(O_t | \pi)} [\text{D}_{\text{KL}}[Q(\phi | O_t, \pi) \parallel Q(\phi | \pi)]]}_{\text{Epistemic value (model parameters information gain)}} \\
 &\quad - \underbrace{\mathbb{E}_{Q(O_t | \pi)} [\ln \tilde{p}(O_t)]}_{\text{Pragmatic value (utility)}} \\
 &\quad + \underbrace{\mathbb{E}_{Q(O_t | \pi)} [\text{D}_{\text{KL}}[Q(D_t, \phi | O_t, \pi) \parallel p(D_t, \phi | O_t, \pi)]]}_{\text{Expected variational approximation error } (\geq 0)}.
 \end{aligned} \tag{A.2}$$

Data availability

The implementation code used for the experiments presented in Section 4 is available in the public repository activeDT Torzoni (2026). The code implements the proposed active digital twin framework and can be used to simulate and generate the plots for digital state estimation, future prediction, and policy inference, as reported in this paper. The observational data used to run the experiments, along with the deep learning models trained according to the implementation details provided in the Appendix of Torzoni et al. (2024), are also available in the same repository. The Matlab library for finite element simulation and reduced-order modeling of partial differential equations employed to generate these data is available in the repository Redbkit Negri (2016).

References

- Achenbach, J.D., 2009. Structural health monitoring — What is the prescription? *Mech. Res. Commun.* 36 (2), 137–142. <http://dx.doi.org/10.1016/j.mechrescom.2008.08.011>.
- Albarracín, M., Hipólito, I., Tremblay, S.E., Fox, J.G., René, G., Friston, K., Ramstead, M.J.D., 2024. Designing explainable artificial intelligence with active inference: A framework for transparent introspection and decision-making. In: Buckley, C.L., et al. (Eds.), *International Workshop on Active Inference*. Cham, Switzerland, pp. 123–144. http://dx.doi.org/10.1007/978-3-031-47958-8_9.
- Aloimonos, J., Weiss, I., Bandyopadhyay, A., 1988. Active vision. *Int. J. Comput. Vis.* 1 (4), 333–356. <http://dx.doi.org/10.1007/BF00133571>.
- Amunts, K., Axer, M., Banerjee, S., Bitsch, L., Bjaalie, J.G., Brauner, P., Brovelli, A., Calarco, N., Carrere, M., Caspers, S., et al., 2024. The coming decade of digital brain research: A vision for neuroscience at the intersection of technology and computing. *Imaging Neurosci.* 2, 1–35. http://dx.doi.org/10.1162/imag_a.00137.
- Andriotis, C.P., Papakonstantinou, K.G., Chatzi, E.N., 2021. Value of structural health information in partially observable stochastic environments. *Struct. Saf.* 93, 102072. <http://dx.doi.org/10.1016/j.strusafe.2020.102072>.
- Arcieri, G., Hoelzl, C., Schwery, O., Straub, D., Papakonstantinou, K., Chatzi, E., 2024. POMDP inference and robust solution via deep reinforcement learning: An application to railway optimal maintenance. *Mach. Learn.* <http://dx.doi.org/10.1007/s10994-024-06559-2>.
- Avci, O., Abdeljaber, O., Kiranyaz, S., Hussein, M., Gabbouj, M., Inman, D., 2021. A review of vibration-based damage detection in civil structures: From traditional methods to machine learning and deep learning applications. *Mech. Syst. Signal Process.* 147, 107077. <http://dx.doi.org/10.1016/j.ymssp.2020.107077>.
- Bajcsy, R., 1988. Active perception. *Proc. IEEE* 76 (8), 966–1005. <http://dx.doi.org/10.1109/5.5968>.
- Bauer, P., Stevens, B., Hazeleger, W., 2021. A digital twin of Earth for the green transition. *Nat. Clim. Chang.* 11 (2), 80–83. <http://dx.doi.org/10.1038/s41558-021-00986-y>.
- Bishop, C.M., 2006. *Pattern Recognition and Machine Learning*. Information Science and Statistics, Springer, New York, New York.
- Bounceur, A., Kara, M., 2025. Intelligent acting digital twins (IADT). *IEEE Access* 13, 15201–15214. <http://dx.doi.org/10.1109/ACCESS.2025.3532545>.
- Buckley, C.L., Kim, C.S., McGregor, S., Seth, A.K., 2017. The free energy principle for action and perception: A mathematical review. *J. Math. Psych.* 81, 55–79. <http://dx.doi.org/10.1016/j.jmp.2017.09.004>.
- Cadena, C., Carlone, L., Carrillo, H., Latif, Y., Scaramuzza, D., Neira, J., Reid, I., Leonard, J.J., 2016. Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Trans. Robot.* 32 (6), 1309–1332. <http://dx.doi.org/10.1109/TRO.2016.2624754>.
- Capellari, G., Chatzi, E., Mariani, S., 2018. Structural health monitoring sensor network optimization through Bayesian experimental design. *ASCE-ASME J. Risk Uncertain. Eng. Syst. Part A: Civ. Eng.* 4 (2), 04018016. <http://dx.doi.org/10.1061/AJRUAE.0000966>.
- Chaudhuri, A., Pash, G., Hormuth, D.A., Lorenzo, G., Kapteyn, M., Wu, C., Lima, E.A.B.F., Yankeelov, T.E., Willcox, K., et al., 2023. Predictive digital twin for optimizing patient-specific radiotherapy regimens under uncertainty in high-grade gliomas. *Front. Artif. Intell.* 6, <http://dx.doi.org/10.3389/frai.2023.1222612>.
- Chinesta, F., Huerta, A., Rozza, G., Willcox, K., 2017. Model reduction methods. In: Stein, E., de Borst, R., Hughes, T.J.R. (Eds.), *Encyclopedia of Computational Mechanics*, Second Edition. John Wiley & Sons, pp. 1–36.
- Chollet, F., et al., 2015. Keras. <https://keras.io>.
- Corral-Acero, J., Margara, F., Marciniak, M., Rodero, C., Loncaric, F., Feng, Y., Gilbert, A., Fernandes, J.F., Bukhari, H.A., Wajdan, A., et al., 2020. The ‘digital twin’ to enable the vision of precision cardiology. *Eur. Heart J.* 41 (48), 4556–4564. <http://dx.doi.org/10.1093/eurheartj/ehaa159>.
- Cotoarbă, D., Straub, D., Smith, I.F.C., 2024. Probabilistic digital twins for geotechnical design and construction. <http://dx.doi.org/10.48550/arXiv.2412.09432>, arXiv preprint arXiv:2412.09432v1.
- Digital Twin: Definition & Value — An AIAA and AIA Position Paper, 2020. AIAA Digital Engineering Integration Committee and others. Tech. Rep., AIAA: Reston, Virginia.
- Eftekhar Azam, S., Mariani, S., 2018. Online damage detection in structural systems via dynamic inverse analysis: A recursive Bayesian approach. *Eng. Struct.* 159, 28–45. <http://dx.doi.org/10.1016/j.engstruct.2017.12.031>.
- Farrar, C.R., Doebling, S.W., Nix, D.A., 2001. Vibration-based structural damage identification. *Philos. Trans. R. Soc. A: Math. Phys. Eng. Sci.* 359 (1778), 131–149. <http://dx.doi.org/10.1098/rsta.2000.0717>.
- Ferrari, A., Willcox, K., 2024. Digital twins in mechanical and aerospace engineering. *Nat. Comput. Sci.* 4, 178–183. <http://dx.doi.org/10.1038/s43588-024-00613-8>.
- Fink, O., Wang, Q., Svensen, M., Dersin, P., Lee, W.-J., Ducoffe, M., 2020. Potential, challenges and future directions for deep learning in prognostics and health management applications. *Eng. Appl. Artif. Intell.* 92, 103678. <http://dx.doi.org/10.1016/j.engappai.2020.103678>.
- FitzGerald, T.H.B., Dolan, R.J., Friston, K., 2015. Dopamine, reward learning, and active inference. *Front. Comput. Neurosci.* 9, 136. <http://dx.doi.org/10.3389/fncom.2015.00136>.
- Fountas, Z., Sajid, N., Mediano, P., Friston, K., 2020. Deep active inference agents using Monte-Carlo methods. In: *Proceedings of the 34th International Conference on Neural Information Processing Systems*. Curran Associates Inc., Red Hook, New York, pp. 11662–11675.
- Friston, K., 2010. The free-energy principle: A unified brain theory? *Nature Rev. Neurosci.* 11 (2), 127–138. <http://dx.doi.org/10.1038/nrn2787>.
- Friston, K.J., Lin, M., Frith, C.D., Pezzulo, G., Hobson, J.A., Ondobaka, S., 2017a. Active inference, curiosity and insight. *Neural Comput.* 29 (10), 2633–2683. http://dx.doi.org/10.1162/neco_a.00999.
- Friston, K.J., Parr, T., de Vries, B., 2017b. The graphical brain: Belief propagation and active inference. *Netw. Neurosci.* 1 (4), 381–414. http://dx.doi.org/10.1162/NETN_a.00018.
- Friston, K., Parr, T., Zeidman, P., 2018. Bayesian model reduction. <http://dx.doi.org/10.48550/arXiv.1805.07092>, arXiv preprint arXiv:1805.07092.
- Friston, K., Rigoli, F., Ognibene, D., Mathys, C., FitzGerald, T., Pezzulo, G., 2015. Active inference and epistemic value. *Cogn. Neurosci.* 6 (4), 187–214. <http://dx.doi.org/10.1080/17588928.2015.1020053>.
- Friston, K., Samothrakis, S., Montague, R., 2012. Active inference and agency: Optimal control without cost functions. *Biol. Cybernet.* 106 (8), 523–541. <http://dx.doi.org/10.1007/s00422-012-0512-8>.
- García-Macias, E., Ubertini, F., 2022. Integrated SHM systems: Damage detection through unsupervised learning and data fusion. In: Cury, A., Ribeiro, D., Ubertini, F., Todd, M.D. (Eds.), *Structural Health Monitoring Based on Data Science Techniques*. Springer International Publishing, Cham, Switzerland, pp. 247–268. http://dx.doi.org/10.1007/978-3-030-81716-9_12.
- Glaser, S.D., Tolman, A., 2008. Sense of sensing: From data to informed decisions for the built environment. *J. Infrastruct. Syst.* 14 (1), 4–14. [http://dx.doi.org/10.1061/\(ASCE\)1076-0342\(2008\)14:1\(4\)](http://dx.doi.org/10.1061/(ASCE)1076-0342(2008)14:1(4)).
- Gregoretto, F., Pezzullo, G., Maisto, D., 2024. cpp-AIF: A multi-core C++ implementation of active inference for partially observable Markov decision processes. *Neurocomputing* 568, 127065. <http://dx.doi.org/10.1016/j.neucom.2023.127065>.
- Grigorescu, S., Trasnea, B., Cocias, T., Macesanu, G., 2020. A survey of deep learning techniques for autonomous driving. *J. Field Robot.* 37 (3), 362–386. <http://dx.doi.org/10.1002/rob.21918>.
- Heins, C., Millidge, B., Da Costa, L., Mann, R.P., Friston, K.J., Couzin, I.D., 2024. Collective behavior from surprise minimization. *Proc. Natl. Acad. Sci.* 121 (17), e2320239121. <http://dx.doi.org/10.1073/pnas.2320239121>.
- Heins, C., Millidge, B., Demekas, D., Klein, B., Friston, K., Couzin, I.D., Tschantz, A., 2022. pymdp: A Python library for active inference in discrete state spaces. *J. Open Source Softw.* 7 (73), 4098. <http://dx.doi.org/10.21105/joss.04098>.
- Henao-García, S., Kapteyn, M., Willcox, K.E., Tezzele, M., Castroviejo-Fernandez, M., Kim, T., Ambrosino, M., Kolmanovsky, I., Basu, H., Jirwankar, P., Sanfelice, R., 2025. Digital-twin-enabled multi-spacecraft on-orbit operations. In: *AIAA SCITECH 2025 Forum*, p. 1432. <http://dx.doi.org/10.2514/6.2025-1432>.
- Henneking, S., Venkat, S., Ghattas, O., 2025. Goal-oriented real-time Bayesian inference for linear autonomous dynamical systems with application to digital twins for tsunami early warning. <http://dx.doi.org/10.48550/arXiv.2501.14911>, arXiv preprint arXiv:2501.14911v1.
- Hughes, T.J.R., 2000. *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*. Dover Civil and Mechanical Engineering, Dover Publications.
- Jans-Singh, M., Leeming, K., Choudhary, R., Girolami, M., 2020. Digital twin of an urban-integrated hydroponic farm. *Data-Centric Eng.* 1, 20. <http://dx.doi.org/10.1017/dce.2020.21>.
- Kaelbling, L.P., Littman, M.L., Cassandra, A.R., 1998. Planning and acting in partially observable stochastic domains. *Artificial Intelligence* 101 (1–2), 99–134. [http://dx.doi.org/10.1016/S0004-3702\(98\)00023-X](http://dx.doi.org/10.1016/S0004-3702(98)00023-X).
- Kamariotis, A., Vlachas, K., Nertimanis, V., Koune, I., Cicirello, A., Chatzi, E., 2024. On the consistent classification and treatment of uncertainties in structural health monitoring applications. *ASCE-ASME J. Risk Uncertain. Eng. Syst. Part B: Mech. Eng.* 11 (1), 011108. <http://dx.doi.org/10.1115/1.4067140>.

- Kapteyn, M.G., Pretorius, J.V.R., Willcox, K.E., 2021. A probabilistic graphical model foundation for enabling predictive digital twins at scale. *Nat. Comput. Sci.* 1 (5), 337–347. <http://dx.doi.org/10.1038/s43588-021-00069-0>.
- Koller, D., Friedman, N., 2009. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, Cambridge, Massachusetts.
- Kritzing, W., Karner, M., Traar, G., Henjes, J., Sihm, W., 2018. Digital twin in manufacturing: A categorical literature review and classification. *IFAC-PapersOnLine* 51 (11), 1016–1022. <http://dx.doi.org/10.1016/j.ifacol.2018.08.474>, 16th IFAC Symposium on Information Control Problems in Manufacturing INCOM 2018.
- Lanillos, P., Meo, C., Pezzato, C., Meera, A.A., Baioumy, M., Ohata, W., Tschantz, A., Millidge, B., Wisse, M., Buckley, C.L., Tani, J., 2021. Active inference in robotics and artificial agents: Survey and challenges. <http://dx.doi.org/10.48550/arXiv.2112.01871>, arXiv preprint [arXiv:2112.01871](https://arxiv.org/abs/2112.01871).
- Li, C., Mahadevan, S., Ling, Y., Choze, S., Wang, L., 2017. Dynamic Bayesian network for aircraft wing health monitoring digital twin. *AIAA J.* 55 (3), 930–941. <http://dx.doi.org/10.2514/1.J055201>.
- Maisto, D., Donnarumma, F., Pezzulo, G., 2024. Interactive inference: A multi-agent model of cooperative joint actions. *IEEE Trans. Syst. Man Cybern.: Syst.* 54 (2), 704–715. <http://dx.doi.org/10.1109/TSMC.2023.3312585>.
- Malings, C., Pozzi, M., 2018. Value-of-information in spatio-temporal systems: Sensor placement and scheduling. *Reliab. Eng. Syst. Saf.* 172, 45–57. <http://dx.doi.org/10.1016/j.res.2017.11.019>.
- Mazzaglia, P., Verbelen, T., Dhoedt, B., 2021. Contrastive active inference. In: *Proceedings of the 35th International Conference on Neural Information Processing Systems*. Curran Associates Inc., Red Hook, New York, pp. 13870–13882.
- McClellan, A., Lorenzetti, J., Pavone, M., Farhat, C., 2022. A physics-based digital twin for model predictive control of autonomous unmanned aerial vehicle landing. *Philos. Trans. R. Soc. A: Math. Phys. Eng. Sci.* 380 (2229), 20210204. <http://dx.doi.org/10.1098/rsta.2021.020417>.
- Murphy, K.P., 2023. *Probabilistic Machine Learning: Advanced Topics*. MIT Press, Cambridge, MA.
- Narouie, V., Wessels, H., Cirak, F., Römer, U., 2025. Mechanical state estimation with a polynomial-chaos-based statistical finite element method. *Comput. Methods Appl. Mech. Engrg.* 441, 117970. <http://dx.doi.org/10.1016/j.cma.2025.117970>.
- National Academy of Engineering and National Academies of Sciences, Engineering, and Medicine, 2024. Foundational research gaps and future directions for digital twins. The National Academies Press, Washington, DC, <http://dx.doi.org/10.17226/26894>.
- Negri, F., 2016. Redbkit, version 2.2. <http://redbkit.github.io/redbKIT>.
- Parr, T., Friston, K.J., 2027. Uncertainty, epistemics and active inference. *J. R. Soc. Interface* 14 (136), 20170376. <http://dx.doi.org/10.1098/rsif.2017.0376>.
- Parr, T., Pezzulo, G., Friston, K.J., 2022. Active Inference: The Free Energy Principle in Mind, Brain, and Behavior. MIT Press, Cambridge, MA, <http://dx.doi.org/10.7551/mitpress/12441.001.0001>.
- Pezzulo, G., Rigoli, F., Friston, K.J., 2018. Hierarchical active inference: A theory of motivated control. *Trends Cogn. Sci.* 22 (4), 294–306. <http://dx.doi.org/10.1016/j.tics.2018.01.009>.
- Phua, A., Davies, C., Delaney, G., 2022. A digital twin hierarchy for metal additive manufacturing. *Comput. Ind.* 140, 103667. <http://dx.doi.org/10.1016/j.compind.2022.103667>.
- Quarteroni, A., Manzoni, A., Negri, F., 2015. *Reduced Basis Methods for Partial Differential Equations: An Introduction*, vol. 92. Springer, Cham, Switzerland, <http://dx.doi.org/10.1007/978-3-319-15431-2>.
- Reis, A.L., Andrade-Compos, A., Matos, P., Henggeler Antunes, C., Lopes, M.A., 2025. An energy and cost efficiency model predictive control framework to optimize water supply systems operation. *Appl. Energy* 384, 125478. <http://dx.doi.org/10.1016/j.apenergy.2025.125478>.
- Rosafalco, L., Torzoni, M., Manzoni, A., Mariani, S., Corigliano, A., 2021. Online structural health monitoring by model order reduction and deep learning algorithms. *Comput. Struct.* 255, 106604. <http://dx.doi.org/10.1016/j.compstruc.2021.106604>.
- Russell, S.J., 2020. *Artificial Intelligence: A Modern Approach*. Pearson Education, London, United Kingdom.
- San, O., Rasheed, A., Bozdemir, E., Deng, J., 2026. The evolution of digital twins from reactive to agentic systems. *Nat. Comput. Sci.* 6 (1), 6–10. <http://dx.doi.org/10.1038/s43588-025-00944-0>.
- Shafto, M., Conroy, M., Doyle, R., Glaessgen, E., Kemp, C., LeMoigne, J., Wang, L., 2020. Technology Area 11: Modeling, Simulation, Information Technology & Processing. Tech. Rep., National Aeronautics and Space Administration, Washington, District of Columbia, <http://dx.doi.org/10.17226/13354>.
- Sirovich, L., 1987. Turbulence and the dynamics of coherent structures. I. Coherent structures. *Quart. Appl. Math.* 45 (3), 561–571. <http://dx.doi.org/10.1090/qam/910462>.
- Smith, R., Friston, K.J., Whyte, C.J., 2022. A step-by-step tutorial on active inference and its application to empirical data. *J. Math. Psych.* 107, 102632. <http://dx.doi.org/10.1016/j.jmp.2021.102632>.
- Taniguchi, T., Murata, S., Suzuki, M., Ognibene, D., Lanillos, P., Ugur, E., Jamone, L., Nakamura, T., Ciria, A., Lara, B., et al., 2023. World models and predictive coding for cognitive and developmental robotics: Frontiers and challenges. *Adv. Robot.* 37 (13), 780–806. <http://dx.doi.org/10.1080/01691864.2023.2225232>.
- Tezzele, M., Carr, S., Topcu, U., Willcox, K.E., 2024. Adaptive planning for risk-aware predictive digital twins. <http://dx.doi.org/10.48550/arXiv.2407.20490>, arXiv preprint [arXiv:2407.20490v2](https://arxiv.org/abs/2407.20490v2).
- Thelen, A., Zhang, X., Fink, O., Lu, Y., Ghosh, S., Youn, B.D., Todd, M.D., Mahadevan, S., Hu, C., Hu, Z., 2022. A comprehensive review of digital twin – part 1: Modeling and twinning enabling technologies. *Struct. Multidiscip. Optim.* 65, 354. <http://dx.doi.org/10.1007/s00158-022-03425-4>.
- Thrun, S., Burgard, W., Fox, D., 2005. *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. MIT Press, Cambridge, MA, <http://dx.doi.org/10.7551/mitpress/12441.001.0001>.
- Torzoni, M., 2024. *Model-Based and Data-Driven Methodologies Toward Predictive Digital Twins of Structures* (Ph.D. thesis). Politecnico di Milano.
- Torzoni, M., 2026. activeDT, GitHub. <https://github.com/MatteoTorzy/activeDT>.
- Torzoni, M., Manzoni, A., Mariani, S., 2022a. Structural health monitoring of civil structures: A diagnostic framework powered by deep metric learning. *Comput. Struct.* 271, 106858. <http://dx.doi.org/10.1016/j.compstruc.2022.106858>.
- Torzoni, M., Manzoni, A., Mariani, S., 2023a. A deep neural network, multi-fidelity surrogate model approach for Bayesian model updating in SHM. In: Rizzo, P., Milazzo, A. (Eds.), *European Workshop on Structural Health Monitoring*. Springer International Publishing, Cham, Switzerland, pp. 1076–1086. http://dx.doi.org/10.1007/978-3-031-07258-1_108.
- Torzoni, M., Manzoni, A., Mariani, S., 2023b. A multi-fidelity surrogate model for structural health monitoring exploiting model order reduction and artificial neural networks. *Mech. Syst. Signal Process.* 197, 110376. <http://dx.doi.org/10.1016/j.ymssp.2023.110376>.
- Torzoni, M., Manzoni, A., Mariani, S., 2025. Enhancing Bayesian model updating in structural health monitoring via learnable mapping. *Eng. Comput.* 41, 4953–4975. <http://dx.doi.org/10.1007/s00366-025-02224-x>.
- Torzoni, M., Rosafalco, L., Manzoni, A., Mariani, S., Corigliano, A., 2022b. SHM under varying environmental conditions: An approach based on model order reduction and deep learning. *Comput. Struct.* 266, 106790. <http://dx.doi.org/10.1016/j.compstruc.2022.106790>.
- Torzoni, M., Tezzele, M., Mariani, S., Manzoni, A., Willcox, K.E., 2024. A digital twin framework for civil engineering structures. *Comput. Methods Appl. Mech. Engrg.* 418, 116584. <http://dx.doi.org/10.1016/j.cma.2023.116584>.
- Tzachor, A., Sabri, S., Richards, C.E., Rajabifard, A., Acuto, M., 2022. Potential and limitations of digital twins to achieve the sustainable development goals. *Nat. Sustain.* 5 (10), 822–829. <http://dx.doi.org/10.1038/s41893-022-00923-7>.
- Ülker-Kaustell, M., 2009. *Some Aspects of the Dynamic Soil-Structure Interaction of a Portal Frame Railway Bridge* (Ph.D. thesis). (October), KTH Royal Institute of Technology.
- Van de Maele, T., Dhoedt, B., Verbelen, T., Pezzulo, G., 2024. A hierarchical active inference model of spatial alternation tasks and the hippocampal-prefrontal circuit. *Nat. Commun.* 15 (1), 9892. <http://dx.doi.org/10.1038/s41467-024-54257-3>.
- Varetti, E., Torzoni, M., Tezzele, M., Manzoni, A., 2026. Adaptive digital twins for predictive decision-making: Online Bayesian learning of transition dynamics. <http://dx.doi.org/10.48550/arXiv.2512.13919>, arXiv preprint [arXiv:2512.13919](https://arxiv.org/abs/2512.13919).
- Vijayaraghavan, P., Queiſer, J.F., Flores, S.V., Tani, J., 2025. Development of compositionality through interactive learning of language and action of robots. *Sci. Robot.* 10 (98), eadp0751. <http://dx.doi.org/10.1126/scirobotics.adp0751>.
- Wainwright, M.J., Jordan, M.I., 2008. Graphical models, exponential families, and variational inference. *Found. Trends Mach. Learn.* 1 (1–2), 1–305. <http://dx.doi.org/10.1561/2200000001>.
- Zakić, B.D., Ryzynski, A., Guo-Hong, C., Jokela, J., 1991. Classification of damage in concrete bridges. *Mater. Struct.* 24, 268–275. <http://dx.doi.org/10.1007/BF02472082>.
- Zhang, Y., Deshmukh, A., Wang, K.-W., 2023. Embodying multifunctional mechano-intelligence in and through phononic metastructures harnessing physical reservoir computing. *Adv. Sci.* 10 (34), 2305074. <http://dx.doi.org/10.1002/adv.202305074>.