




POLITECNICO
MILANO 1863



**Funded by
the European Union**

Project: 101140720 - IMMENSE - ERC-2023-ADG



European Research Council
Established by the European Commission

DECISION-MAKING AND LEARNING IN MECHANICAL SYSTEMS VIA FREE ENERGY MINIMIZATION

Scuola Superiore Meridionale, June 23th, 2026

Matteo Torzoni¹

Joint work with D. Maisto², A. Manzoni³, F. Donnarumma², G. Pezzulo², and A. Corigliano¹

¹Department of Civil and Environmental Engineering, Politecnico di Milano


²Institute of Cognitive Sciences and Technologies, National Research Council (Roma)

³MOX – Department of Mathematics, Politecnico di Milano

Adaptation and learning in sentient structures

From reactive to autonomous systems, with active data assimilation and learning capabilities (closed perception-action loop).

Dynamic, adaptive, bidirectional interaction



Digital to physical:
Actuation and
Observation

Physical environment

Physical to digital:
Data assimilation
and interpretation

Digital agent




Personalized, dynamically updating,
predictive decision-making (what-if scenarios)

Requirements:


1. **Reaction and learning** without simulations.
2. **Continual adaptation** to novelty.
3. **Minimal computational cost** for low-power devices.


Partially observable Markov decision process (POMDP)

Application agnostic model of physical-virtual interaction for decision-making under uncertainty.



Active inference for POMDPs






POMDP under active inference: State inference (preception)

The **free energy** objective quantifies how much beliefs must change to explain observations:

- **For state inference**, free energy minimization implies updating the digital state.

Conditional independence structure
admits belief state factorization:

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



POMDP under active inference: Policy inference (action)

The **free energy** objective quantifies how much beliefs must change to explain observations:

- **For state inference**, free energy minimization implies updating the digital state.
- **For policy inference:**
 1. **Goal-directed behavior** (utility).
Actions that **realize objectives**.
 2. **Active sensing** (epistemic value).
Actions that **reduce state uncertainty**.
 3. **Curiosity about cause-effect** (novelty).
Actions that **reduce model uncertainty**.

Conditional independence structure
admits belief state factorization:


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

Demonstration of active learning from experiments

Analogous to a structure with tunable parts to infer boundary conditions or stiffness.

Generative process:


- Aluminum bar 1 meter long.
- Variable supports (left, double, right).
- Sliding mass of 1kg (the agent).
- 3 monitored degrees of freedom.




Agent:

1. Slides the mass **to learn** the observation model – initialized as uniform (curiosity).
2. Slides the mass **to understand** the constraint condition (state info gain).
3. Slides the mass **to achieve a mechanical goal**, such as maximizing observed displacements (utility).

Perception-action cycle simulation







Perception-action cycle simulation



- **Spontaneous transition** from exploration to exploitation.
- **Emergent behavior** stabilizing the mass at locations that help disambiguate the context and maximize utility.


Performance comparison and learned observation model






Different observations and mechanical goal

Strain maximization




Strain minimization




Increased model resources through state expansion

Maximize displacements




Digital twin perspective: Two-ways information flow

“Personalized information construct that mimics specific attributes of a natural/engineered system/process. It is updated via data assimilation, has predictive capabilities, and guides actions that maximize value”.



Active digital twins: beyond passive data assimilation





Structural health monitoring workflow and hierarchical structure

Genoa 2018, partial collapse of the Morandi Bridge – 43 fatalities




Aloha Airlines 1988, separation of the fuselage upper lobe – 1 fatality





Why SHM?
More informed assessment of safety, better allocation of resources, increased availability.

Simulation-based damage identification




Given system parameters, compare solution trajectories to sensor recordings

$$\begin{cases} \mathbf{M}_r \ddot{\hat{\mathbf{x}}}(t) + \mathbf{C}_r(\mu) \dot{\hat{\mathbf{x}}}(t) + \mathbf{K}_r(\mu) \hat{\mathbf{x}}(t) = \mathbf{f}_r(t, \mu) \\ \hat{\mathbf{x}}(0) = \mathbf{W}^\top \mathbf{x}_0 \\ \dot{\hat{\mathbf{x}}}(0) = \mathbf{W}^\top \dot{\mathbf{x}}_0 \end{cases} \quad \begin{aligned} \mathbf{M}_r &\equiv \mathbf{W}^\top \mathbf{M} \mathbf{W}, \\ \mathbf{C}_r(\mu) &\equiv \mathbf{W}^\top \mathbf{C}(\mu) \mathbf{W}, \\ \mathbf{K}_r(\mu) &\equiv \mathbf{W}^\top \mathbf{K}(\mu) \mathbf{W}, \\ \mathbf{f}_r(t, \mu) &\equiv \mathbf{W}^\top \mathbf{f}(t, \mu) \end{aligned}$$



Application to structural health monitoring



- Test: 17,292 DOFs
- Train: 133 POMs
- 10 noisy sensors at 400 Hz

Parameters:

- Damage location: $\{0, \Omega_1, \Omega_2, \Omega_3, \Omega_4, \Omega_5, \Omega_6\}$
- Damage level: $[30\%, 80\%]$
- Convoy velocity: $[160, 215]$ km/h
- Axle mass: $[16, 22]$ ton

Generative process: evolving structural health, stochastic degradation (depends on the loads).

Agent:

1. Actively infers the structural health state.
2. Optimizes costs for structural health and actions.
3. Updates model parameters (learning).

Generative model

Digital state factors:

1. Damage location.
2. Damage level.
3. Epistemic switch.


Observations:

1. Pre-classified signals (DL models).


Controls:

1. Do nothing – DN.
 2. Maintenance – MA.
 3. Restrict operations – RO.
 4. Read sensors – RE.
- } Affect the environment
 → Triggers the epistemic state with an informative likelihood


Active digital twin resolves uncertainty through action



Purely goal-directed:
Simulation fails in 50% of the trials.



Goal-directed and information-seeking:
Dynamically adjusted data acquisition
prevents desynchronization.



Thank you for your attention



POLITECNICO
MILANO 1863




**Funded by
the European Union**



European Research Council
Established by the European Commission

Project: 101140720 - IMMENSE - ERC-2023-ADG



Contacts

Matteo Torzoni
Department of Civil and Environmental Engineering

matteo.torzoni@polimi.it
www.erc-immense.eu