Thermodynamics-informed Neural Networks

P. Urdeitx, C. Bermejo, A. Tierz, M. Martinez, L. Tesán, I. Alfaro, D. González, F. Chinesta, E. Cueto

ESI Group-UZ Chair of the National Strategy on AI

España | digital 2026

UKACM-SEMNI Autumn School

ARTS ET MÉTIERS ParisTech

Universidad Zaragoza

Motivation

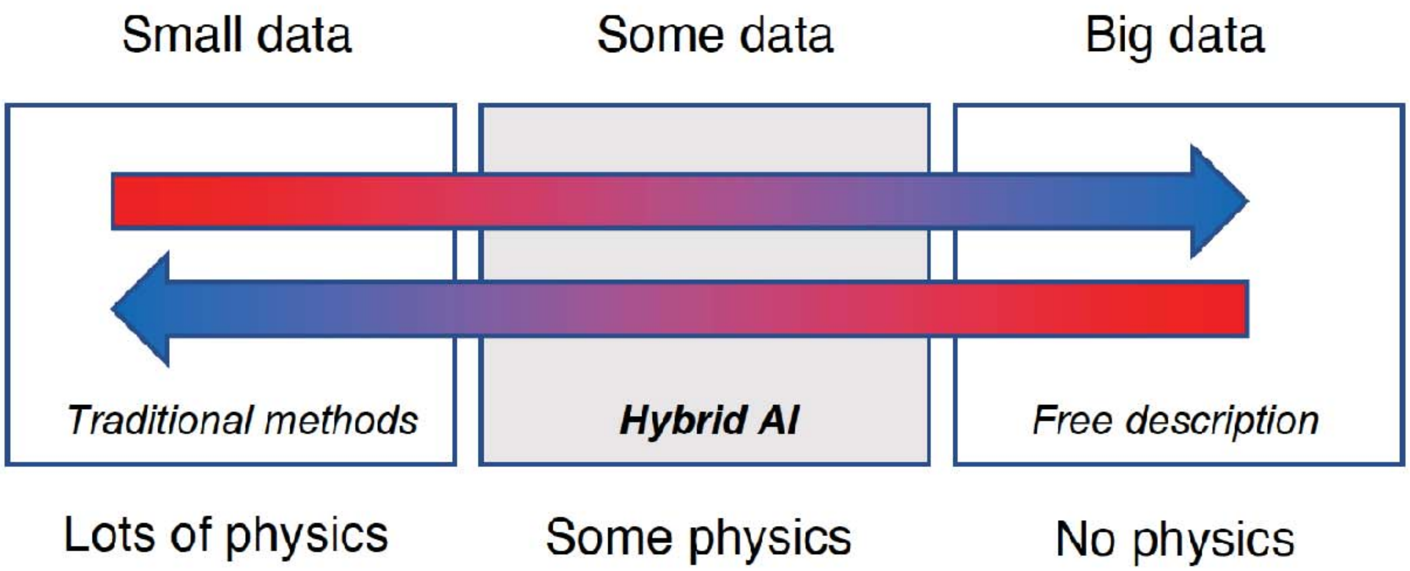

Adapted from G. Karniadakis.

Motivation: world models

"Scientists […] need to build AI that doesn't just operate by matching patterns but can also reason about the physical world". [1]

"It's about modeling the world…" [2]

"… to create machines that can learn internal models of how the world works […], plan how to accomplish complex tasks, and readily adapt to unfamiliar situations." [3]

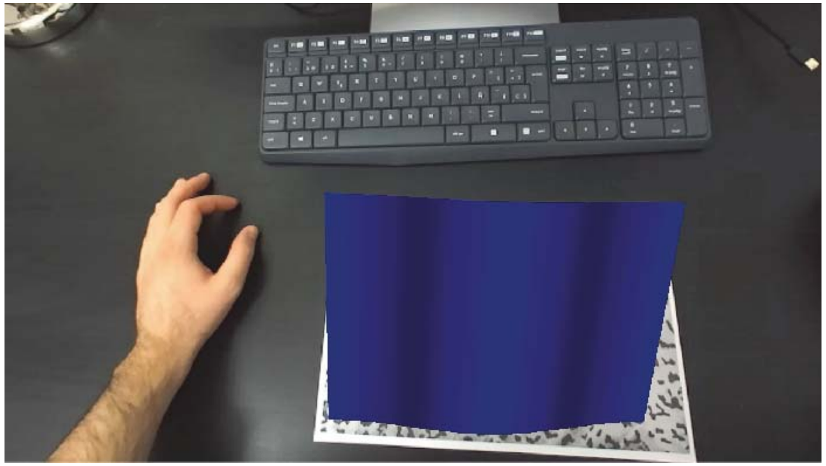

Badias, Alberto, et al. "Morph-DSLAM: Model order reduction for physics-based deformable SLAM." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44.11 (2021): 7764-7777.

[1]. Matthew Hutson, Nature Index, November 17th 2023.
[2]. Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and brain sciences*, *40*, e253.
[3] LeCun, Y. (2022). A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. *Open Review*, *62*.

Cognitive digital twins

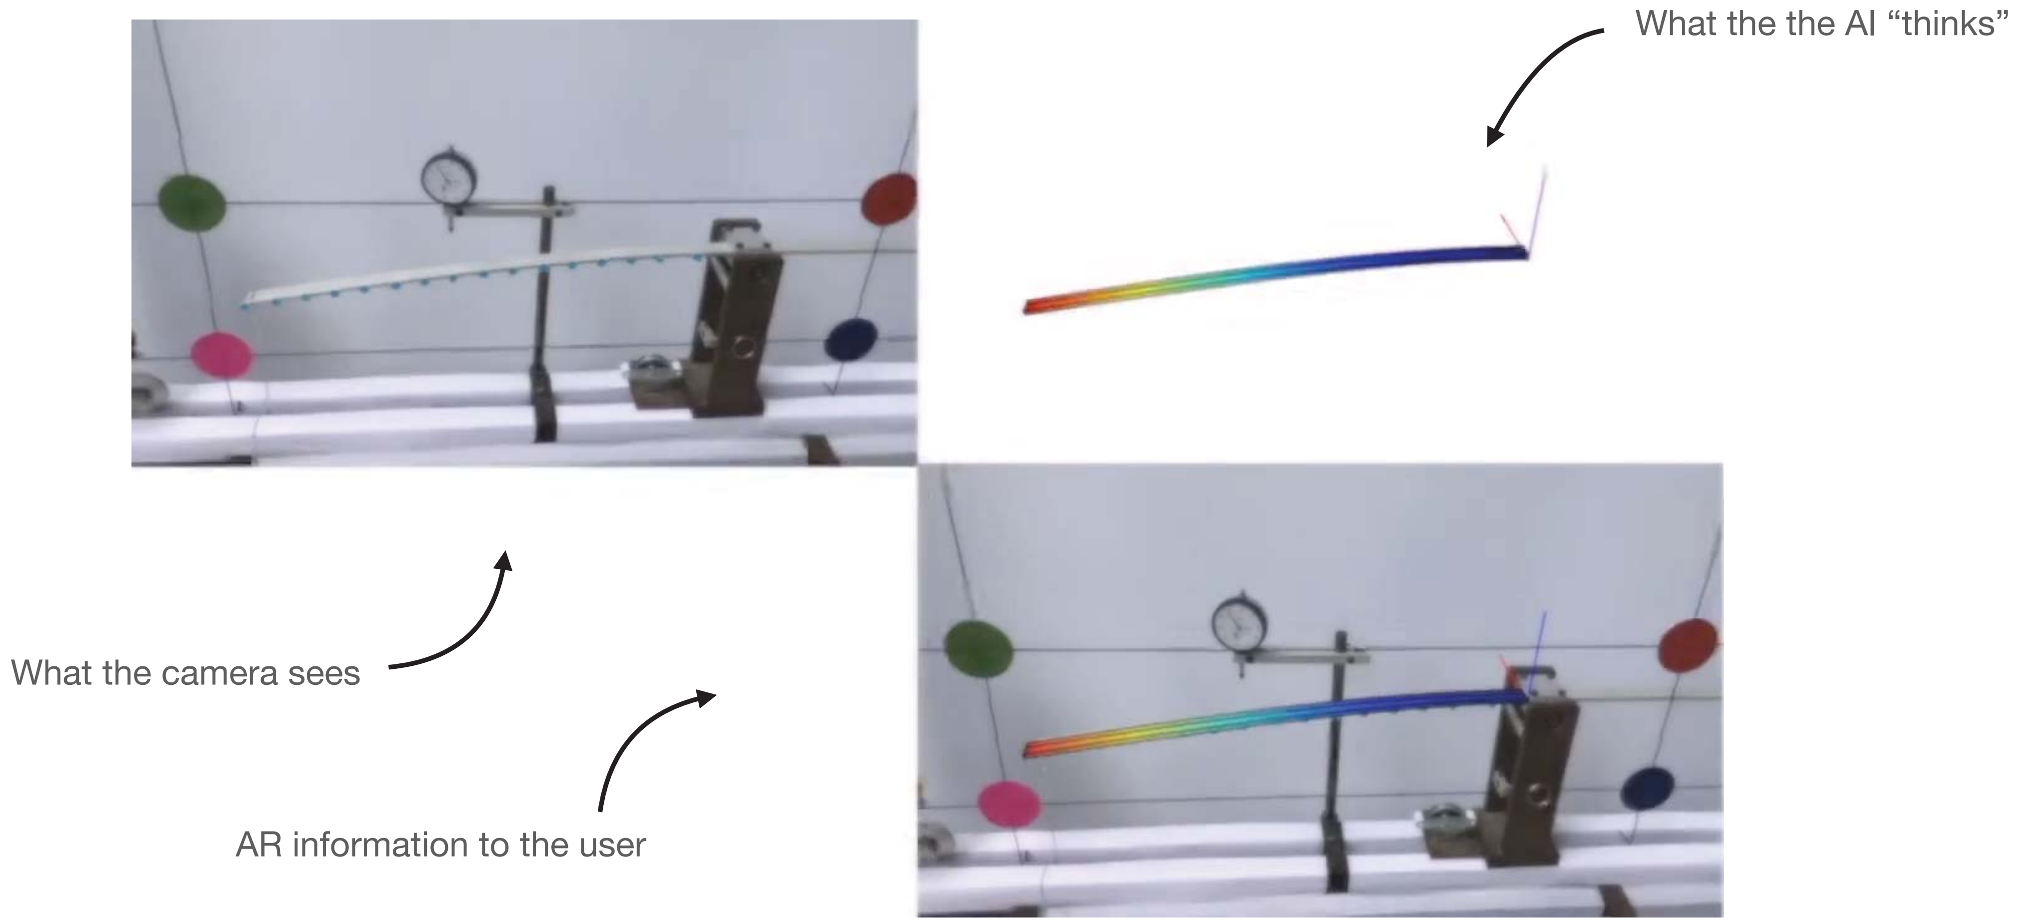

Our approach to cognitive twins

- Able to see (through computer vision)
- Able to understand what they see (perception, machine learning)
- Able to make prognosis (reasoning, real-time simulation)
- Able to inform for decision making (Augmented Reality)

Physics-enhanced machine learning

- Encompasses Scientific ML, Informed ML, Physics-enhanced AI, …
- Provides a framework for guiding high-consequence decision making in engineering applications
- Hybrid physics-data models integrating
  - advanced computational models
  - multi-fidelity data
  - domain knowledge; prior knowledge
  - first principles and appropriate biases

Physics-Enhanced Machine Learning: a position paper for dynamical systems investigations. Alice Cicirello. Arxiv: 2405.05987, 2024.

The importance of inductive biases

Loss landscape

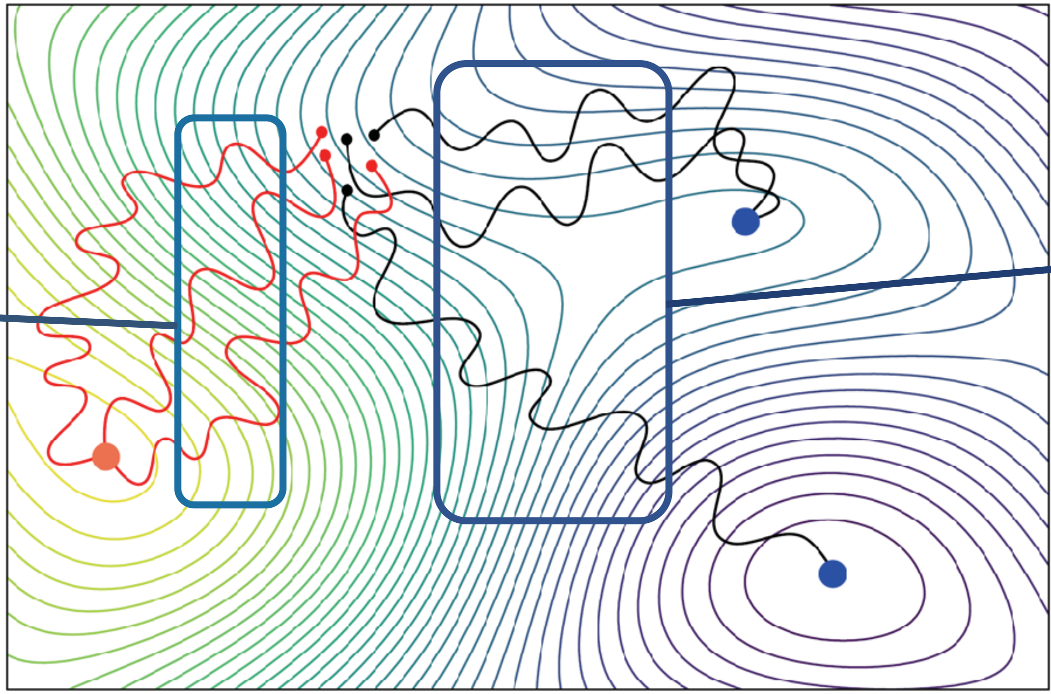

Strong bias
Physics
+
Domain

Weak bias

Taxonomy of biases

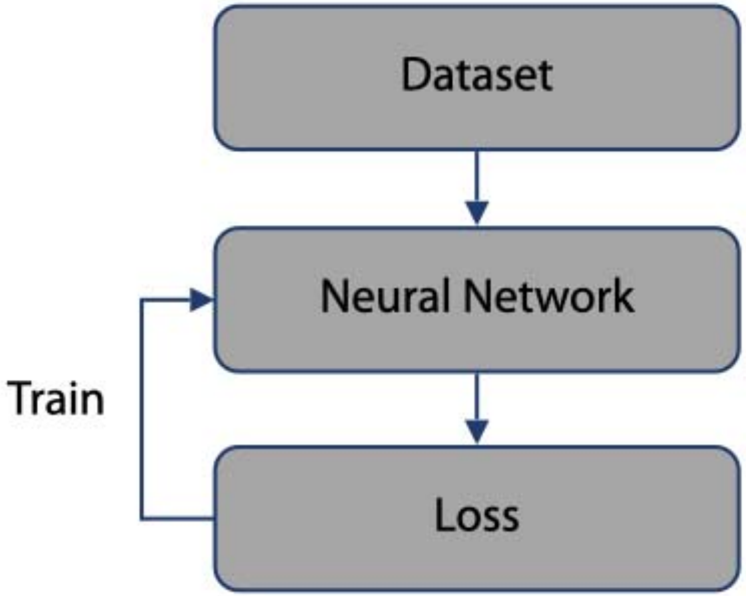

Observational bias

Inductive bias

- NeuralODE [Chen, 2018]
- FNO [Li, 2020]
- MP-GNN [Gilmer, 2017]
- MeshGraphNet [Pfaff, 2020]

Learning bias

- PINNs [Raissi, 2019]
- DeepONet [Lu, 2019]

Contents

The problem of learning physics from data

- Learn a dynamical system from data
- State vector: $\boldsymbol{z} = \left(z_1, z_2, \ldots\right)$

$$\dot{\boldsymbol{z}} = \frac{d\boldsymbol{z}}{dt} = \boxed{F(\boldsymbol{z}, t)}$$

- Time interval: $t \in (0,T]$
- Initial conditions: $\boldsymbol{z}(t=0) = \boldsymbol{z}_0$

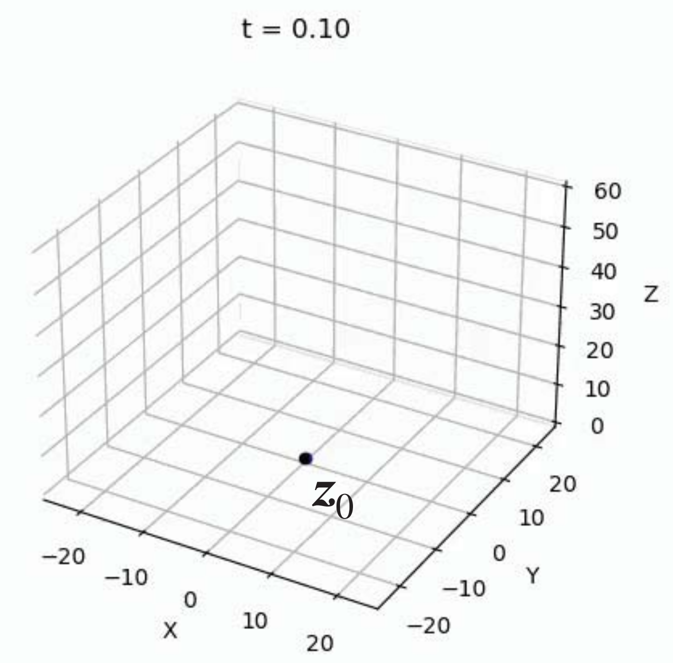

Biases: conservative systems

- Hamiltonian mechanics
  - State variables: $\boldsymbol{z} = (\boldsymbol{q}, \boldsymbol{p})$
  - Hamiltonian: $\mathcal{H} = \mathcal{H}(\boldsymbol{q}, \boldsymbol{p}) = T(\boldsymbol{p}) + V(\boldsymbol{q})$

Hamilton's equations

$$\begin{cases} \dfrac{d\boldsymbol{p}}{dt} = -\dfrac{\partial \mathcal{H}}{\partial \boldsymbol{q}} \\[2ex] \dfrac{d\boldsymbol{q}}{dt} = \dfrac{\partial \mathcal{H}}{\partial \boldsymbol{p}} \end{cases} \rightarrow \begin{pmatrix} \dfrac{dq}{dt} \\[1ex] \dfrac{dp}{dt} \end{pmatrix} = \underbrace{\begin{pmatrix} 0 & \mathbf{1} \\ -\mathbf{1} & 0 \end{pmatrix}}_{\boldsymbol{L}(\boldsymbol{z})} \begin{pmatrix} \dfrac{\partial \mathcal{H}}{\partial q} \\[1ex] \dfrac{\partial \mathcal{H}}{\partial p} \end{pmatrix}$$

- Poisson matrix
- Skew-symm

$$\frac{d\boldsymbol{z}}{dt} = \boldsymbol{L}\frac{\partial E}{\partial \boldsymbol{z}}$$

- Symplectic
- Reversible

Biases: conservative systems

- Hamiltonian mechanics
  - State variables: $\boldsymbol{z} = (\boldsymbol{q}, \boldsymbol{p})$
  - Hamiltonian: $\mathscr{H} = \mathscr{H}(\boldsymbol{q}, \boldsymbol{p}) = T(\boldsymbol{p}) + V(\boldsymbol{q})$

Hamilton's equations

$$\begin{cases} \dfrac{d\boldsymbol{p}}{dt} = -\dfrac{\partial \mathscr{H}}{\partial \boldsymbol{q}} \\[2ex] \dfrac{d\boldsymbol{q}}{dt} = \dfrac{\partial \mathscr{H}}{\partial \boldsymbol{p}} \end{cases}$$

Hamiltonian NN [Sánchez-González, 2019]
SympNets [Jin, 2020]
Lagrangian NN [Bhatoo, 2021]
Poisson NN [Jin, 2023]

$$\frac{d\boldsymbol{z}}{dt} = \boldsymbol{L}\frac{\partial E}{\partial \boldsymbol{z}}$$

- Symplectic
- Reversible

Biases: dissipative systems

- Introduction of a new potential: **entropy**, S
- **Metriplectic** (metric+symplectic) formulation

$$\dot{\boldsymbol{z}}_t = \underbrace{\boldsymbol{L}(\boldsymbol{z}_t)\nabla E(\boldsymbol{z}_t)}_{\text{reversible}} + \underbrace{\boldsymbol{M}(\boldsymbol{z}_t)\nabla S(\boldsymbol{z}_t)}_{\text{irreversible}}, \quad \boldsymbol{z}(0) = \boldsymbol{z}_0.$$

- Bracket structure

$$\frac{d\boldsymbol{z}}{dt} = \{\boldsymbol{z}, E\} + [\boldsymbol{z}, S].$$

GENERIC (Öttinger & Grmela)

- Degeneracy conditions:

$$\boldsymbol{L}(\boldsymbol{z})\nabla S(\boldsymbol{z}) = \boldsymbol{0},$$
$$\boldsymbol{M}(\boldsymbol{z})\nabla E(\boldsymbol{z}) = \boldsymbol{0}.$$

- By choosing $\boldsymbol{L}$ skew-symmetric and $\boldsymbol{M}$ symmetric, positive semi-definite,

$$\dot{E}(\boldsymbol{z}) = \nabla E(\boldsymbol{z}) \cdot \dot{\boldsymbol{z}} = \nabla E(\boldsymbol{z}) \cdot \boldsymbol{L}(\boldsymbol{z})\nabla E(\boldsymbol{z}) + \nabla E(\boldsymbol{z}) \cdot \boldsymbol{M}(\boldsymbol{z})\nabla S(\boldsymbol{z}) = 0,$$

(conservation of energy in closed systems.)

- Equivalently,

$$\dot{S}(\boldsymbol{z}) = \nabla s(\boldsymbol{z}) \cdot \dot{\boldsymbol{z}} = \nabla S(\boldsymbol{z}) \cdot \boldsymbol{L}(\boldsymbol{z})\nabla E(\boldsymbol{z}) + \nabla S(\boldsymbol{z}) \cdot \boldsymbol{M}(\boldsymbol{z})\nabla S(\boldsymbol{z}) \geq 0,$$

(second principle of thermodynamics.)

Structure-preserving neural networks

- Parametrization of GENERIC operators:

$$\boldsymbol{L} = \boldsymbol{l} - \boldsymbol{l}^\top, \qquad \boldsymbol{M} = \boldsymbol{m}\boldsymbol{m}^\top.$$

- Data loss:

$$\mathcal{L}_n^{\text{data}} = \left\| \frac{dz^{\text{GT}}}{dt} - \frac{dz^{\text{net}}}{dt} \right\|_2^2,$$

- Degeneracy loss:

$$\mathcal{L}_n^{\text{deg}} = \left\| \boldsymbol{L} \frac{\partial S}{\partial z_n} \right\|_2^2 + \left\| \boldsymbol{M} \frac{\partial E}{\partial z_n} \right\|_2^2.$$

- Global loss:

$$\mathcal{L} = \frac{1}{N_{\text{batch}}} \sum_{n=0}^{N_{\text{batch}}} (\lambda \mathcal{L}_n^{\text{data}} + \mathcal{L}_n^{\text{deg}}).$$

- Hernández, Q., Badías, A., González, D., Chinesta, F., & Cueto, E. (2021). Structure-preserving neural networks. *Journal of Computational Physics*, *426*, 109950. SOFT
- Lee, K., Trask, N., & Stinis, P. (2021). Machine learning structure preserving brackets for forecasting irreversible processes. *Advances in Neural Information Processing Systems*, *34*, 5696-5707. HARD
- Zhang, Z., Shin, Y., & Em Karniadakis, G. (2022). GFINNs: GENERIC formalism informed neural networks for deterministic and stochastic dynamical systems. *Philosophical Transactions of the Royal Society A*, *380*(2229), 20210207. HARD

Structure-preserving ROMs

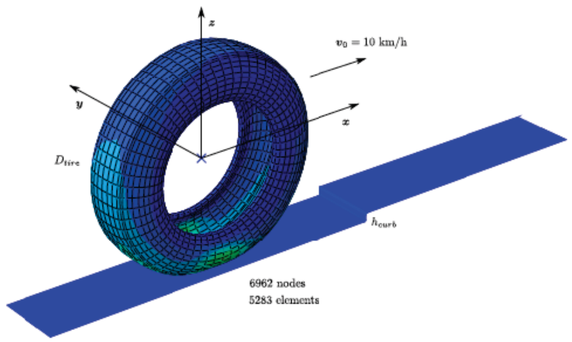

Unveil the intrinsic dimensionality of the manifold—L1 autoencoder

Structure-preserving time integration

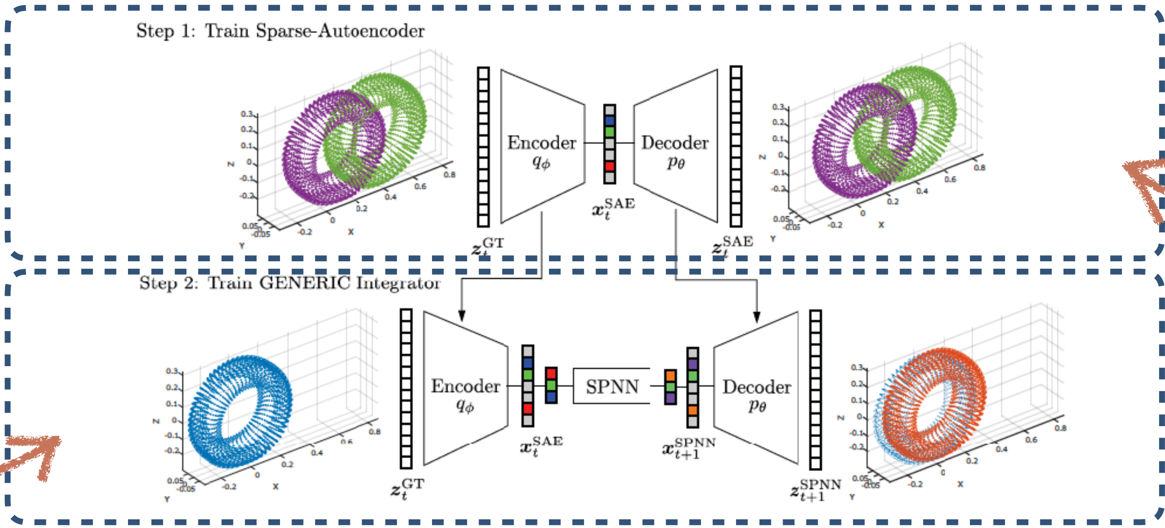

Results

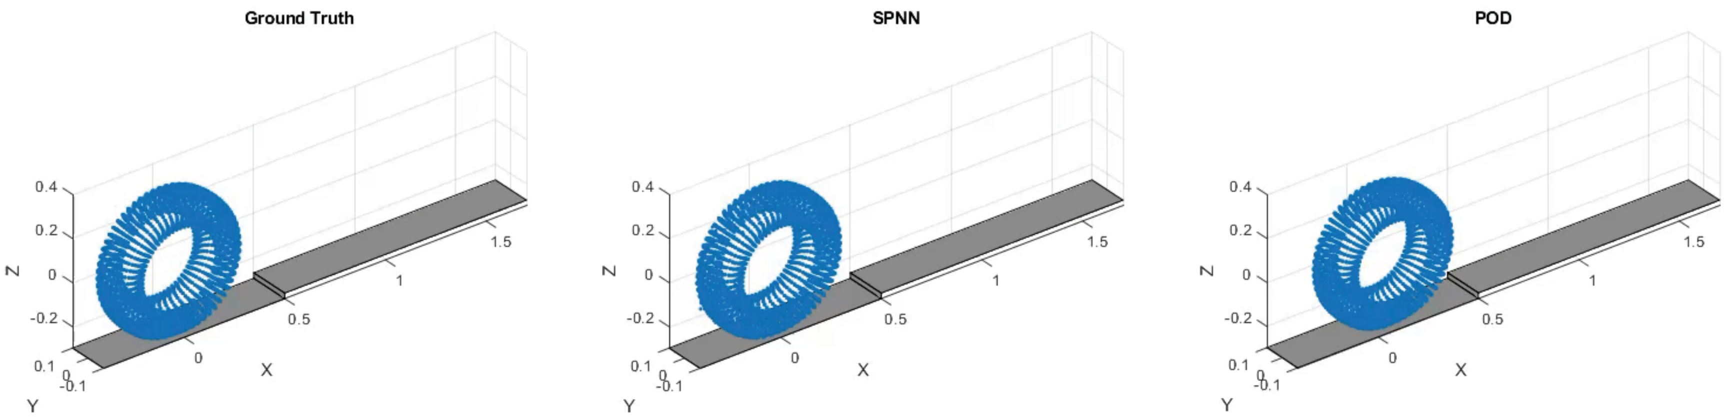

Hernandez, Q., Badias, A., Gonzalez, D., Chinesta, F., & Cueto, E. (2021). Deep learning of thermodynamics-aware reduced-order models from data. *Computer Methods in Applied Mechanics and Engineering*, *379*, 113763.

Take-home message

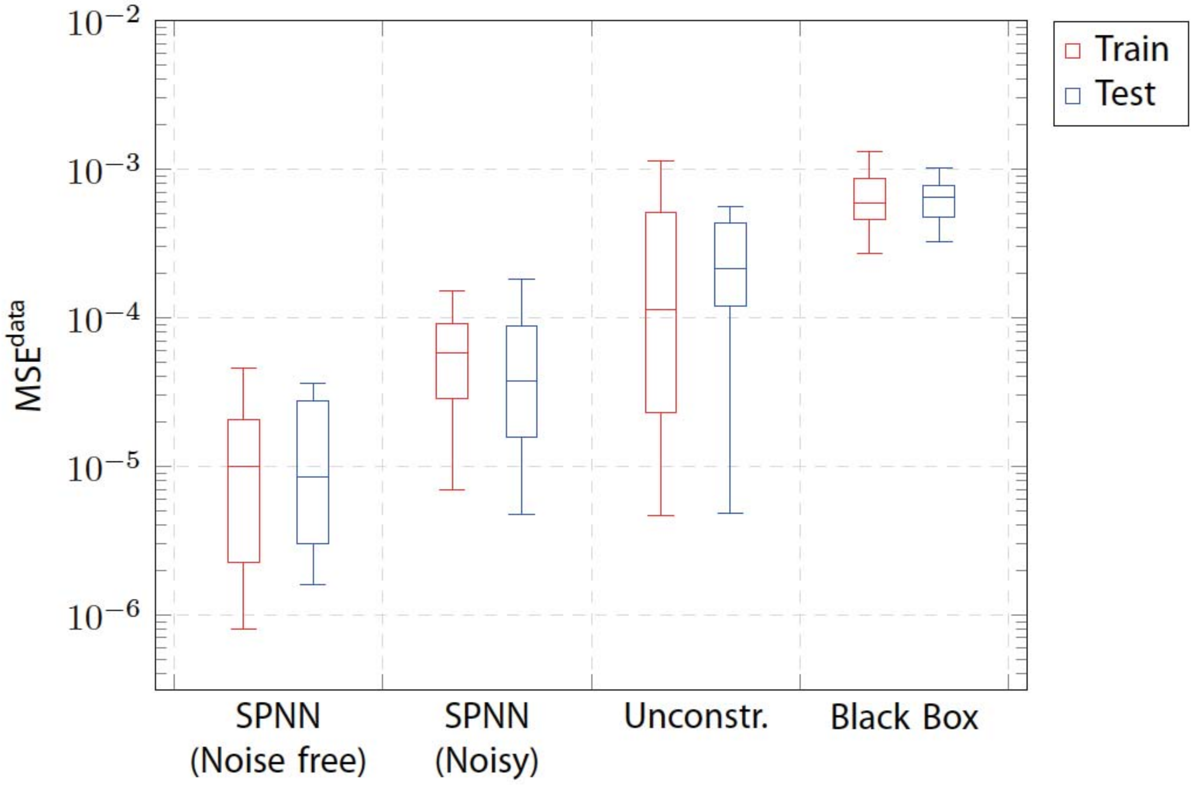

Hernández, Q., Badías, A., González, D., Chinesta, F., & Cueto, E. (2021). Structure-preserving neural networks. *Journal of Computational Physics*, *426*, 109950.

Contents

Geometric bias: Graph Neural Networks

- Graph construction

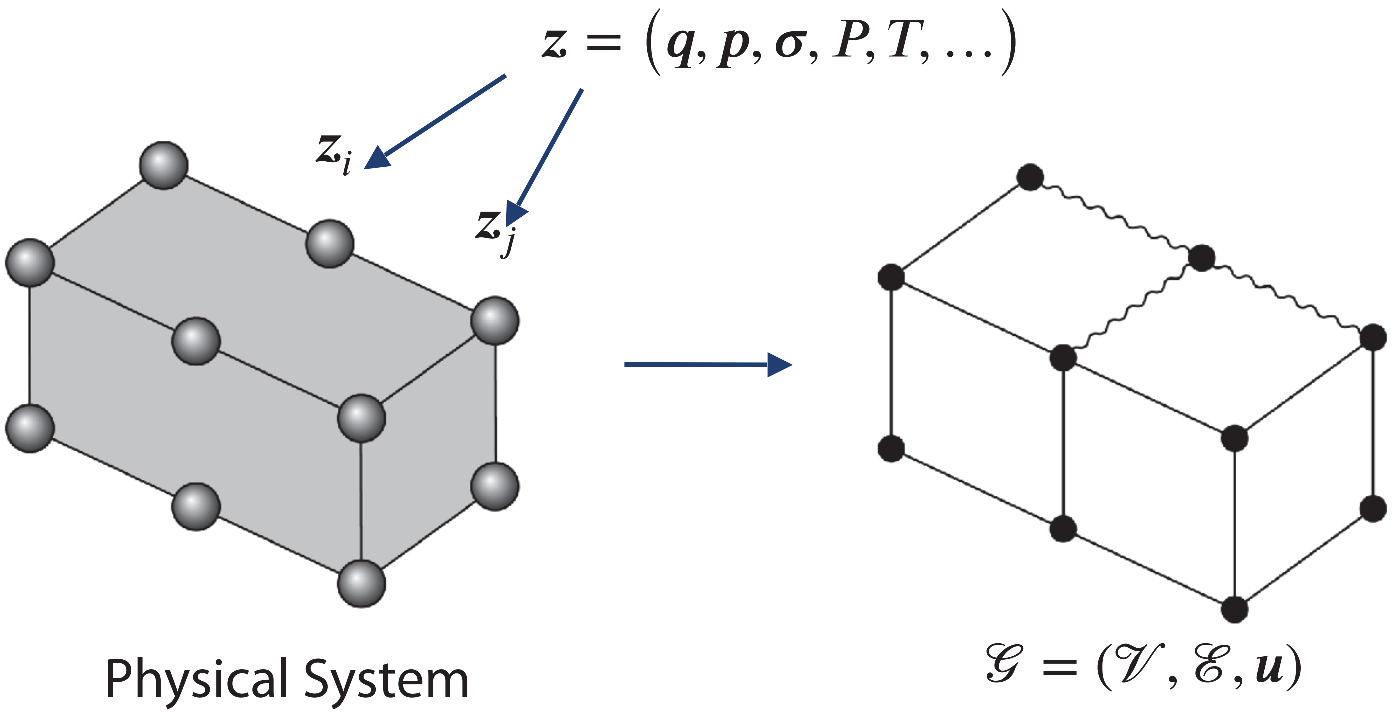

Geometric bias

- Graph construction

$$e_{ij} = \left(q_i - q_j, |q_i - q_j|\right)$$

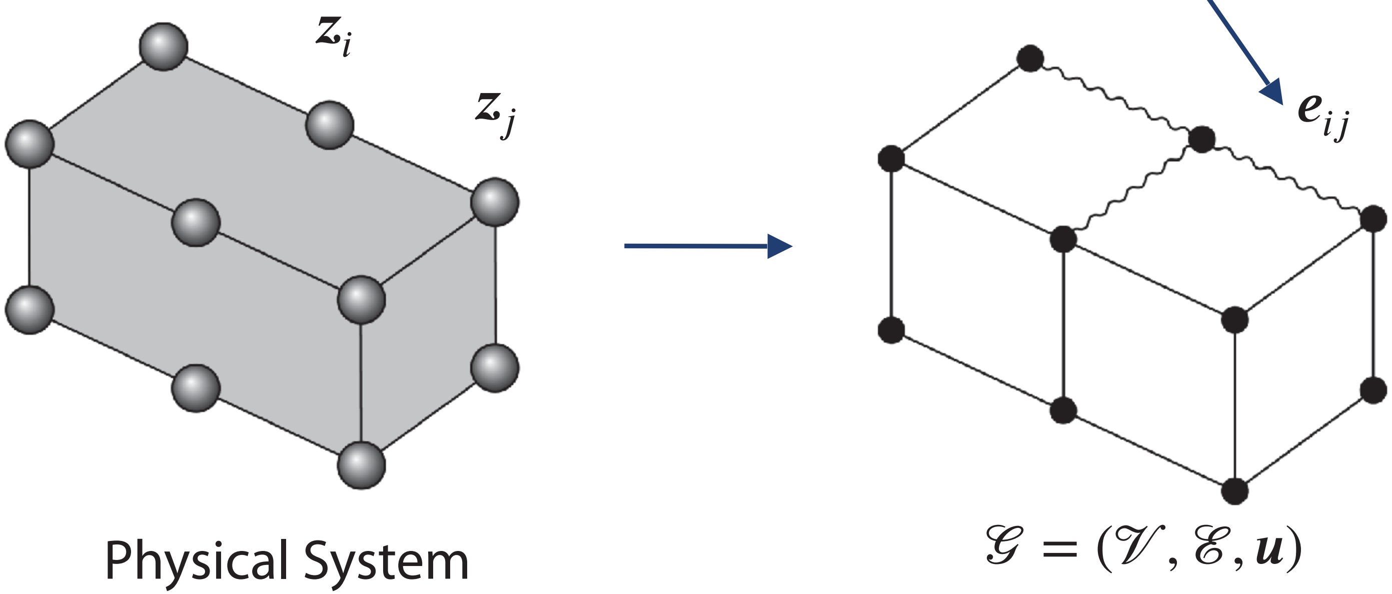

$$\mathcal{G} = (\mathcal{V}, \mathcal{E}, \boldsymbol{u})$$

Translation Equivariant

Geometric bias

- Graph construction

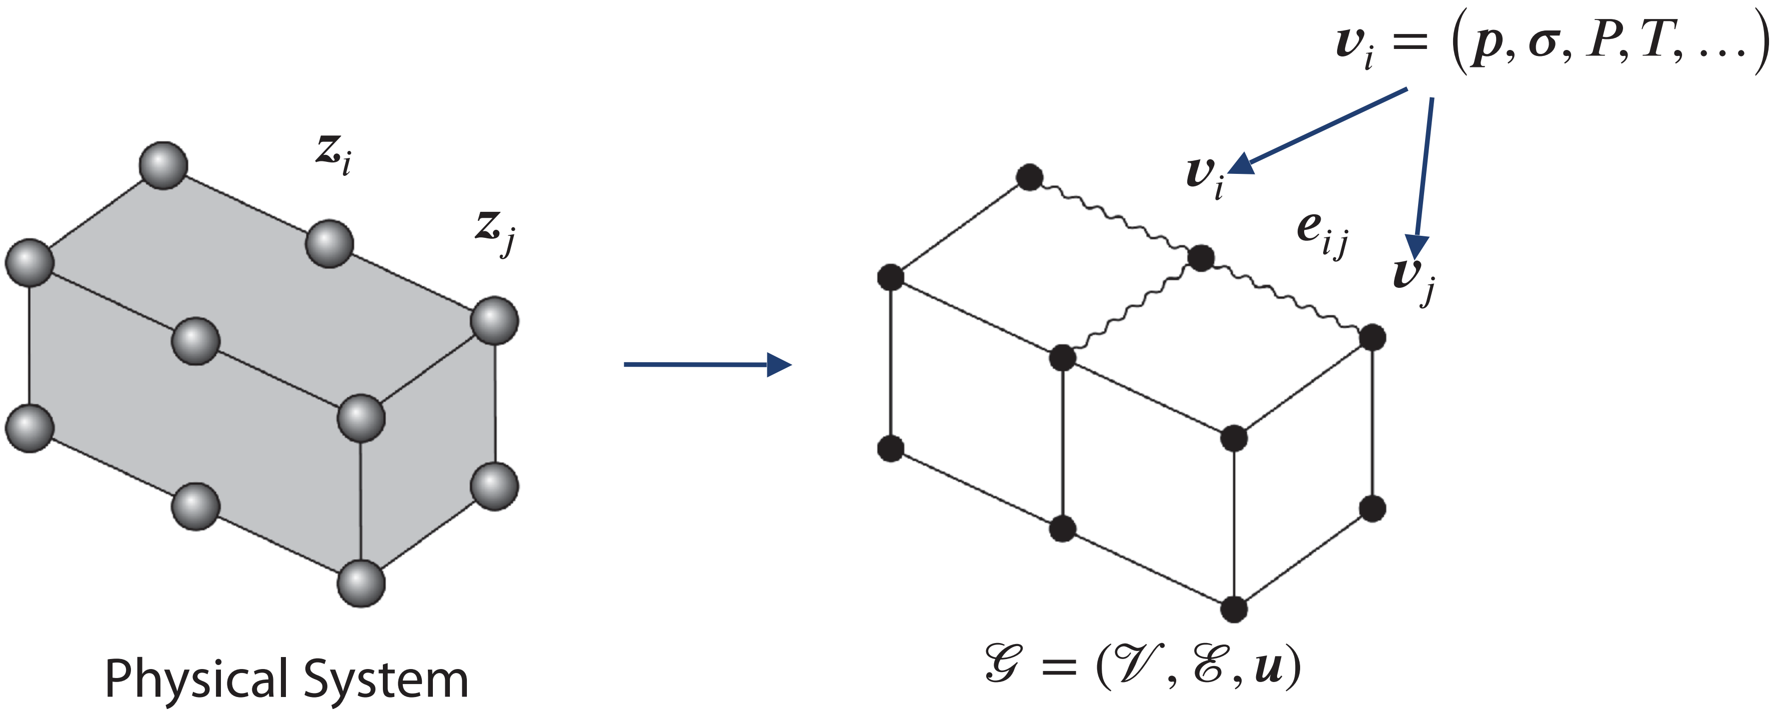

Geometric bias

- Graph construction

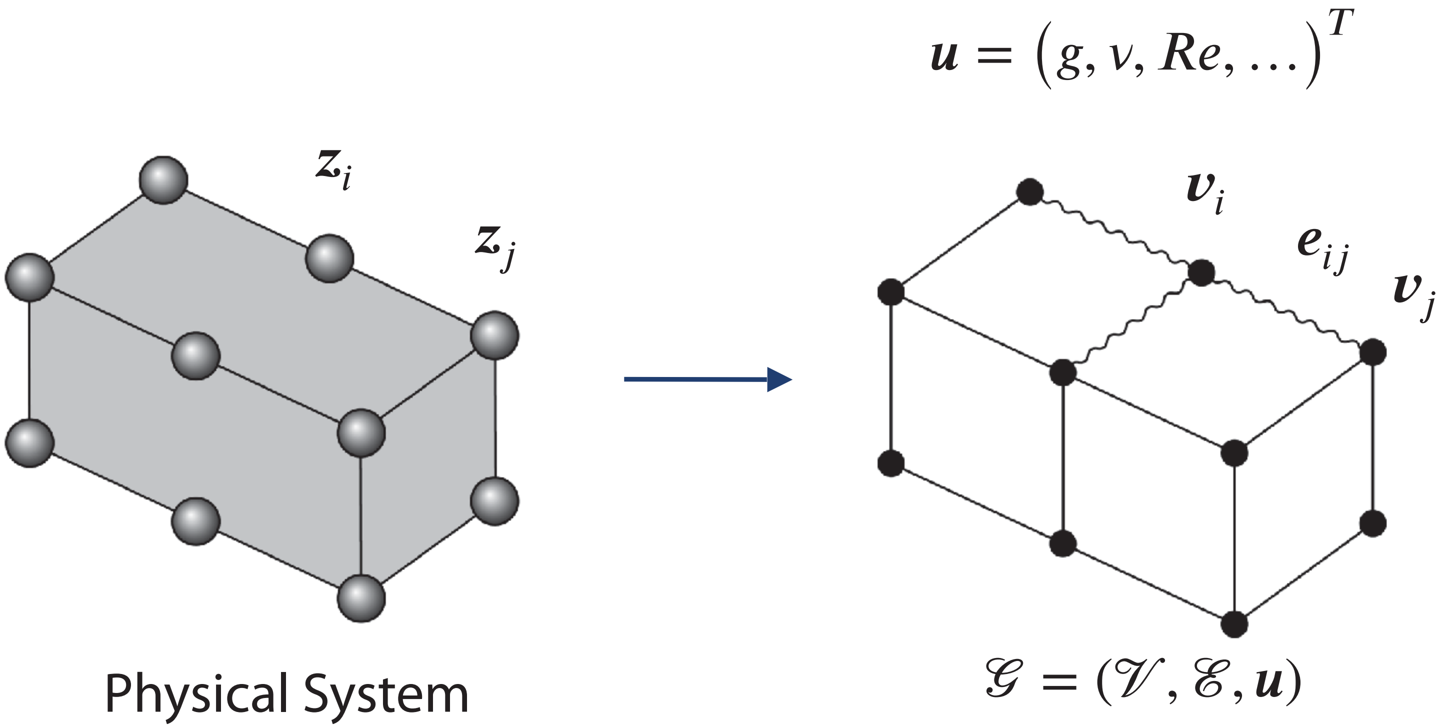

Geometric bias

- Encode – Process – Decode [Battaglia, 2018]

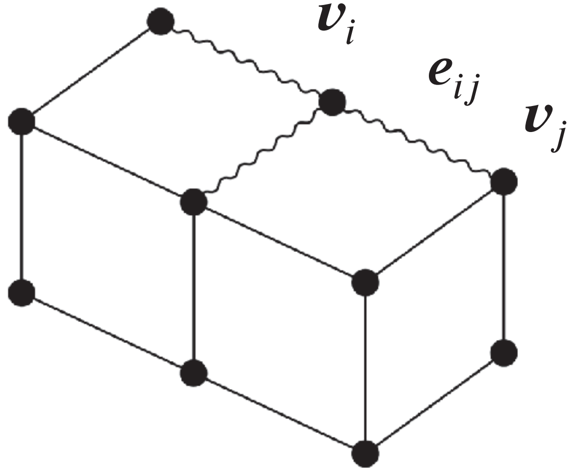

Geometric bias

- Encode – Process – Decode [Battaglia, 2018]
  1. Encoder: ε_v, ε_e

$$\boldsymbol{x}_i = \varepsilon_v(\boldsymbol{v}_i)$$

$$\boldsymbol{x}_{ij} = \varepsilon_e(\boldsymbol{e}_{ij})$$

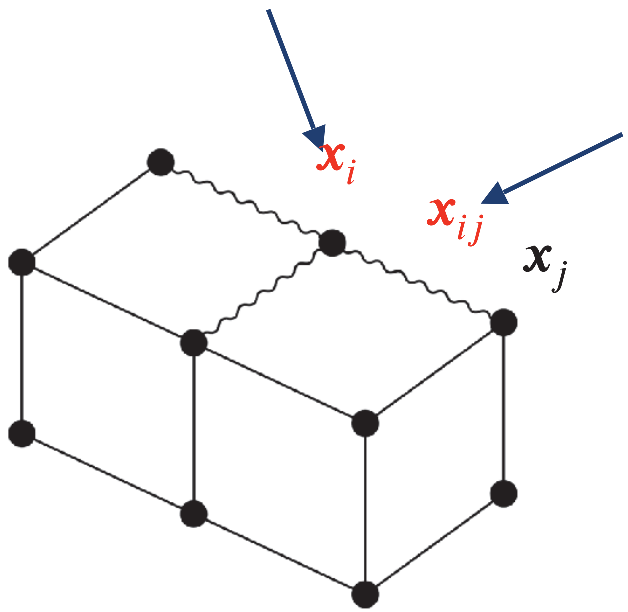

Geometric bias

- Encode – Process – Decode [Battaglia, 2018]
  1. Encoder: $\varepsilon_v, \varepsilon_e$
  2. Update Edges: π_e

$$\boldsymbol{x}'_{ij} = \pi_e(\boldsymbol{x}_{ij}, \boldsymbol{x}_i, \boldsymbol{x}_j, \boldsymbol{u})$$

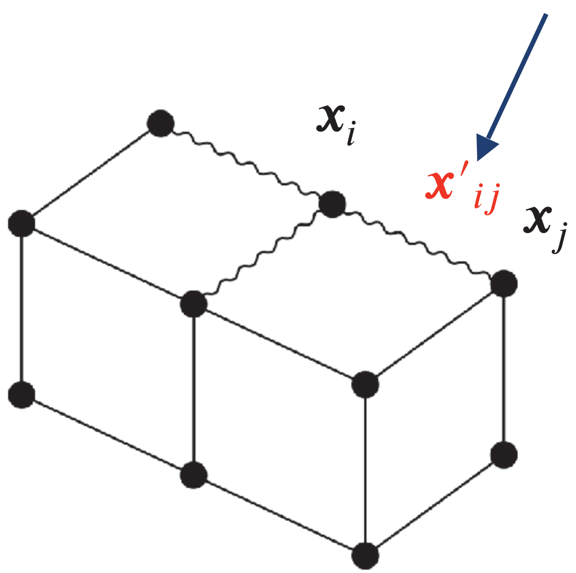

Geometric bias

- Encode – Process – Decode [Battaglia, 2018]
  1. Encoder: $\varepsilon_v, \varepsilon_e$
  2. Update Edges: π_e
  3. Message Passing

$$\boldsymbol{m}_i = \sum_{j \in \mathcal{N}_i} \boldsymbol{x}'_{ij}$$

Permutation Equivariant

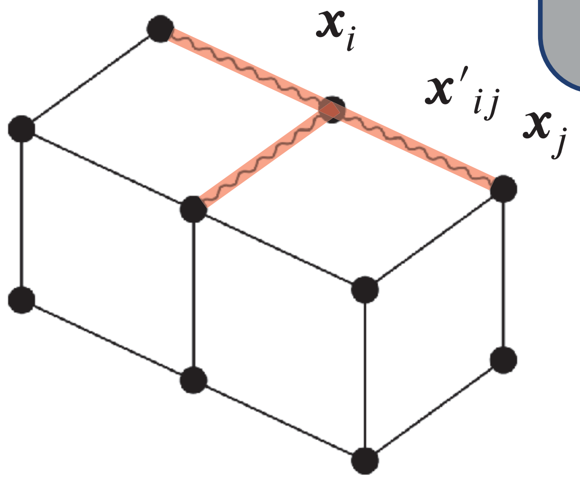

Geometric bias

- Encode – Process – Decode

[Battaglia, 2018]

1. Encoder: $\varepsilon_v, \varepsilon_e$
2. Update Edges: π_e
3. Message Passing
4. Update Vertices: π_v

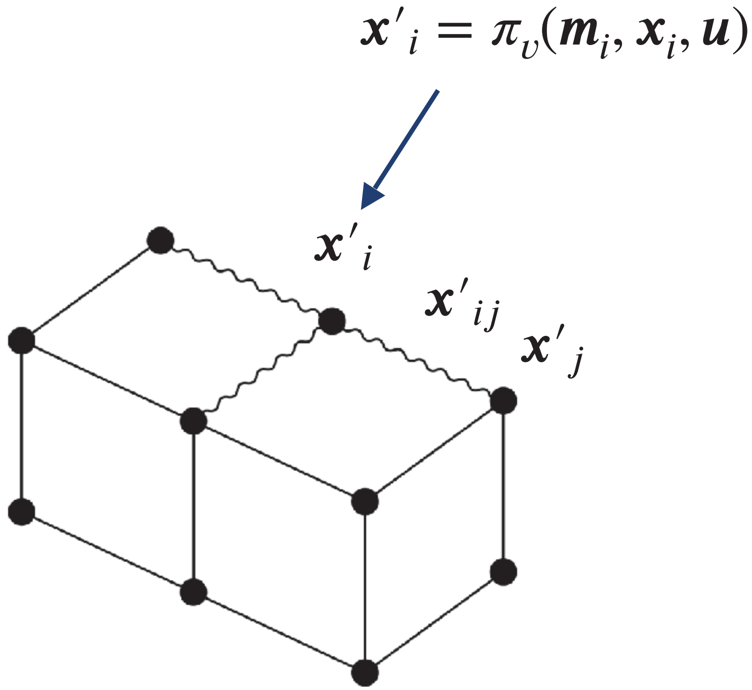

Geometric bias

- Encode – Process – Decode

[Battaglia, 2018]

1. Encoder: $\varepsilon_v, \varepsilon_e$
2. Update Edges: π_e
3. Message Passing
4. Update Vertices: π_v
5. Decoder: δ

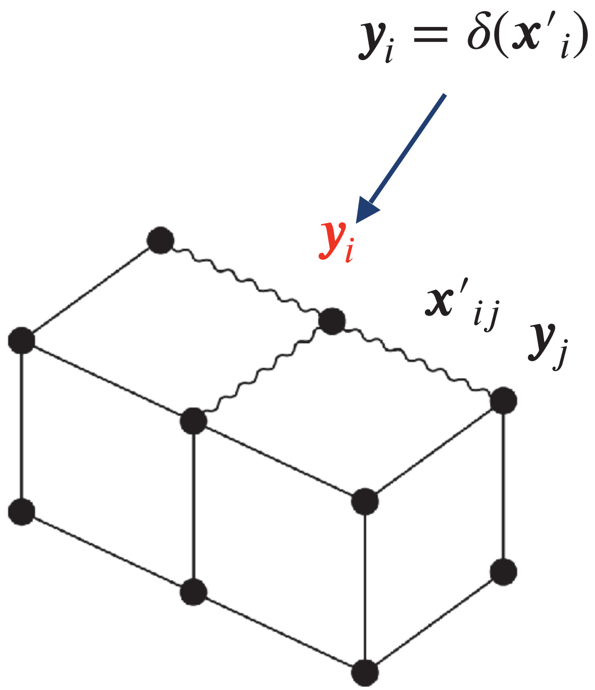

Experiments

- Ablation study

Method	Physics	Geometry
SPNN [Hernández, 2021]	✓	✗
GNN [Pfaff, 2021]	✗	✓
TIGNN [Hernández, 2022]	✓	✓

Experiments

- Bending viscoelastic beam
  - State Space: $\mathcal{S} = \{\boldsymbol{z} = (\boldsymbol{q}, \boldsymbol{v}, \boldsymbol{\sigma})\}$
  - Dataset: 52 load positions

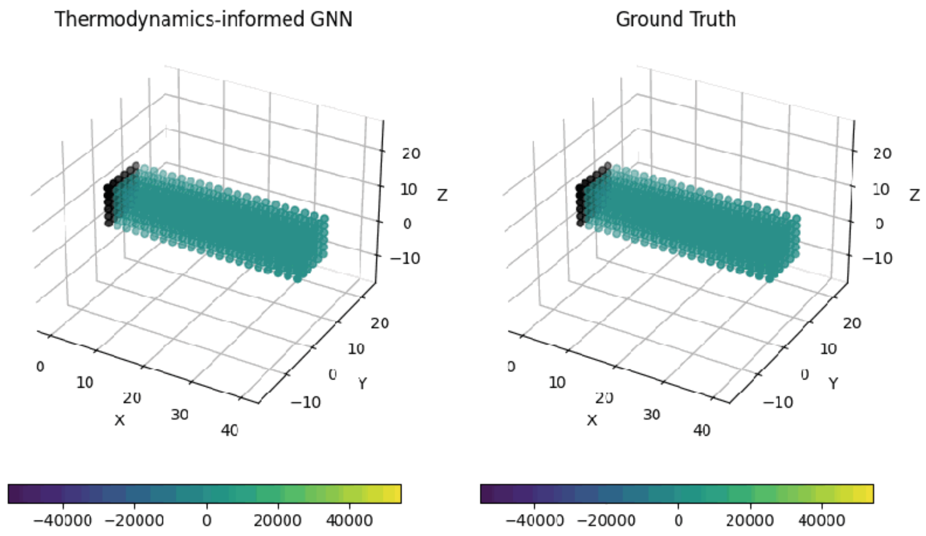

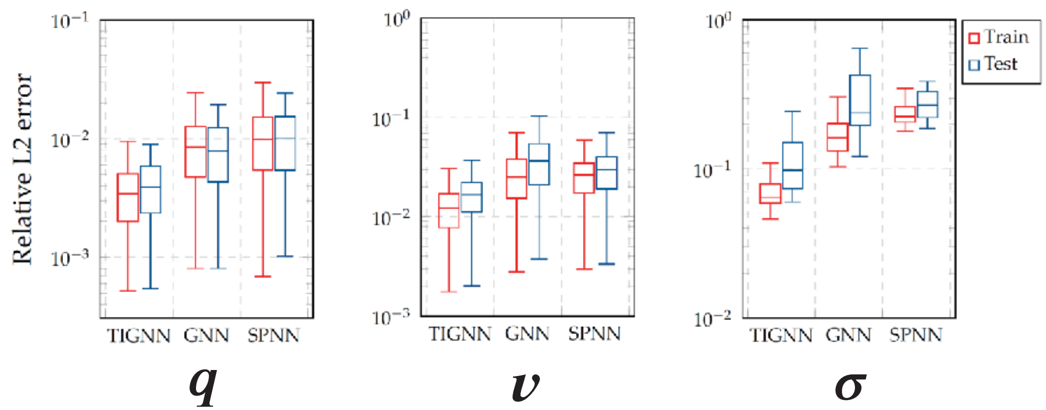

Experiments

- Bending viscoelastic beam
  - State Space: $\mathcal{S} = \{ \boldsymbol{z} = (\boldsymbol{q}, \boldsymbol{v}, \boldsymbol{\sigma}) \}$
  - Dataset: 52 load positions

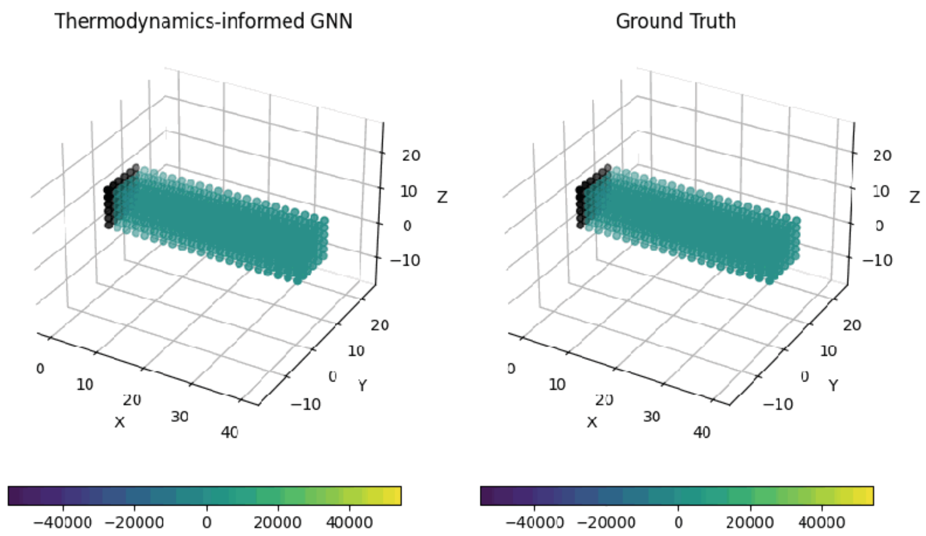

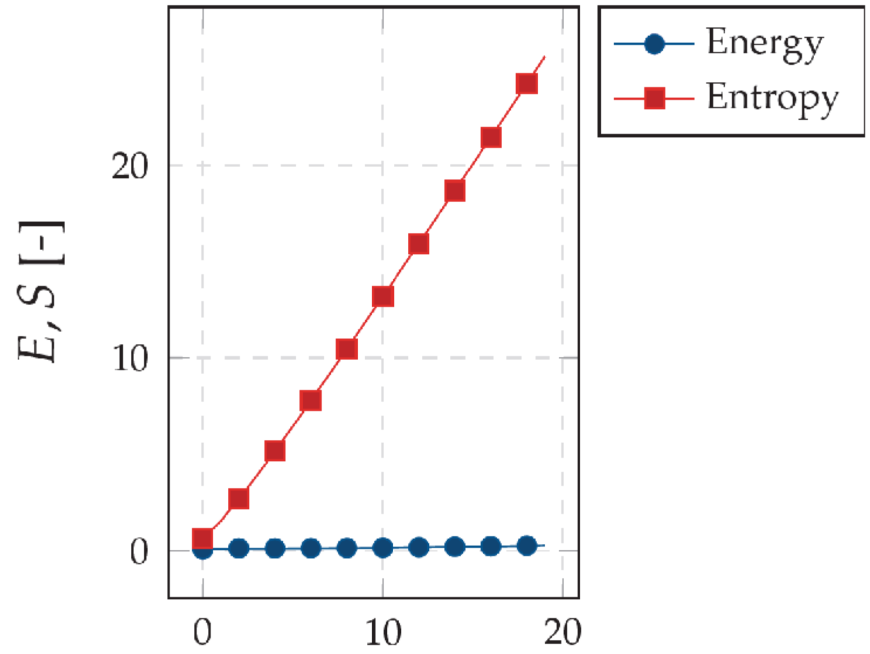

Experiments: previously unseen meshes

- Flow past a cylinder
  - State Space: $\mathcal{S} = \{\boldsymbol{z} = (\boldsymbol{v}, P)\}$
  - Dataset: 30 geometries + $\boldsymbol{v}$

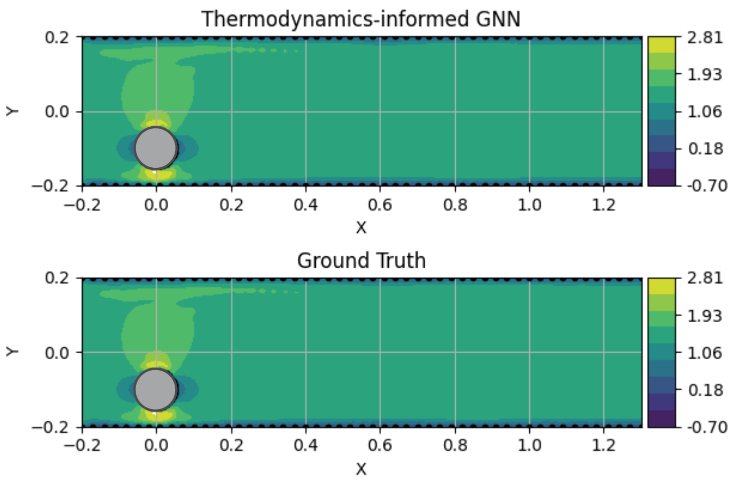

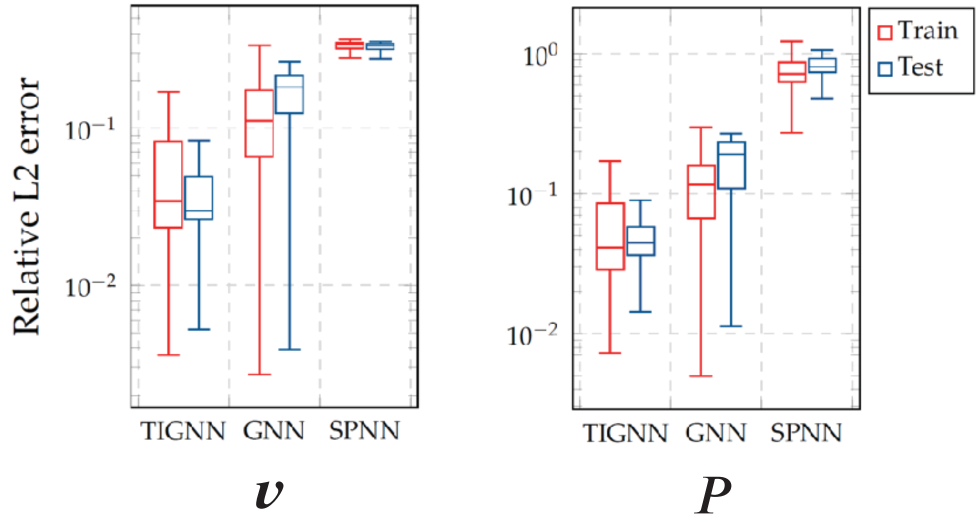

Hernandez, Quercus, et al. "Thermodynamics-informed Graph Neural Networks." *IEEE Transactions on Artificial Intelligence* (2022).

Are GNNs learning the physics?

Training:

Test:

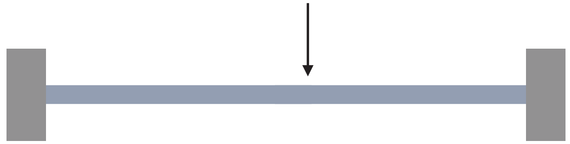

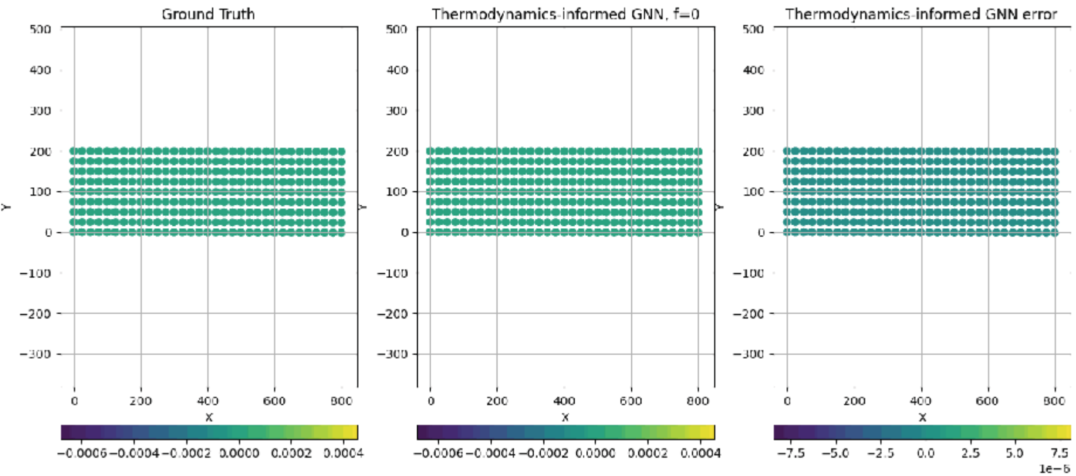

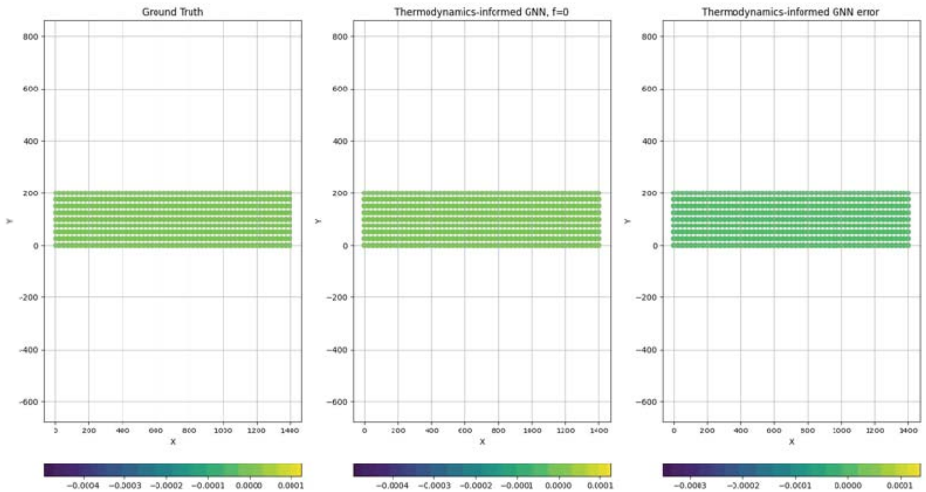

Tierz, Alicia, et al. "Graph neural networks informed locally by thermodynamics." *arXiv preprint arXiv:2405.13093* (2024).

Are GNNs learning the physics?

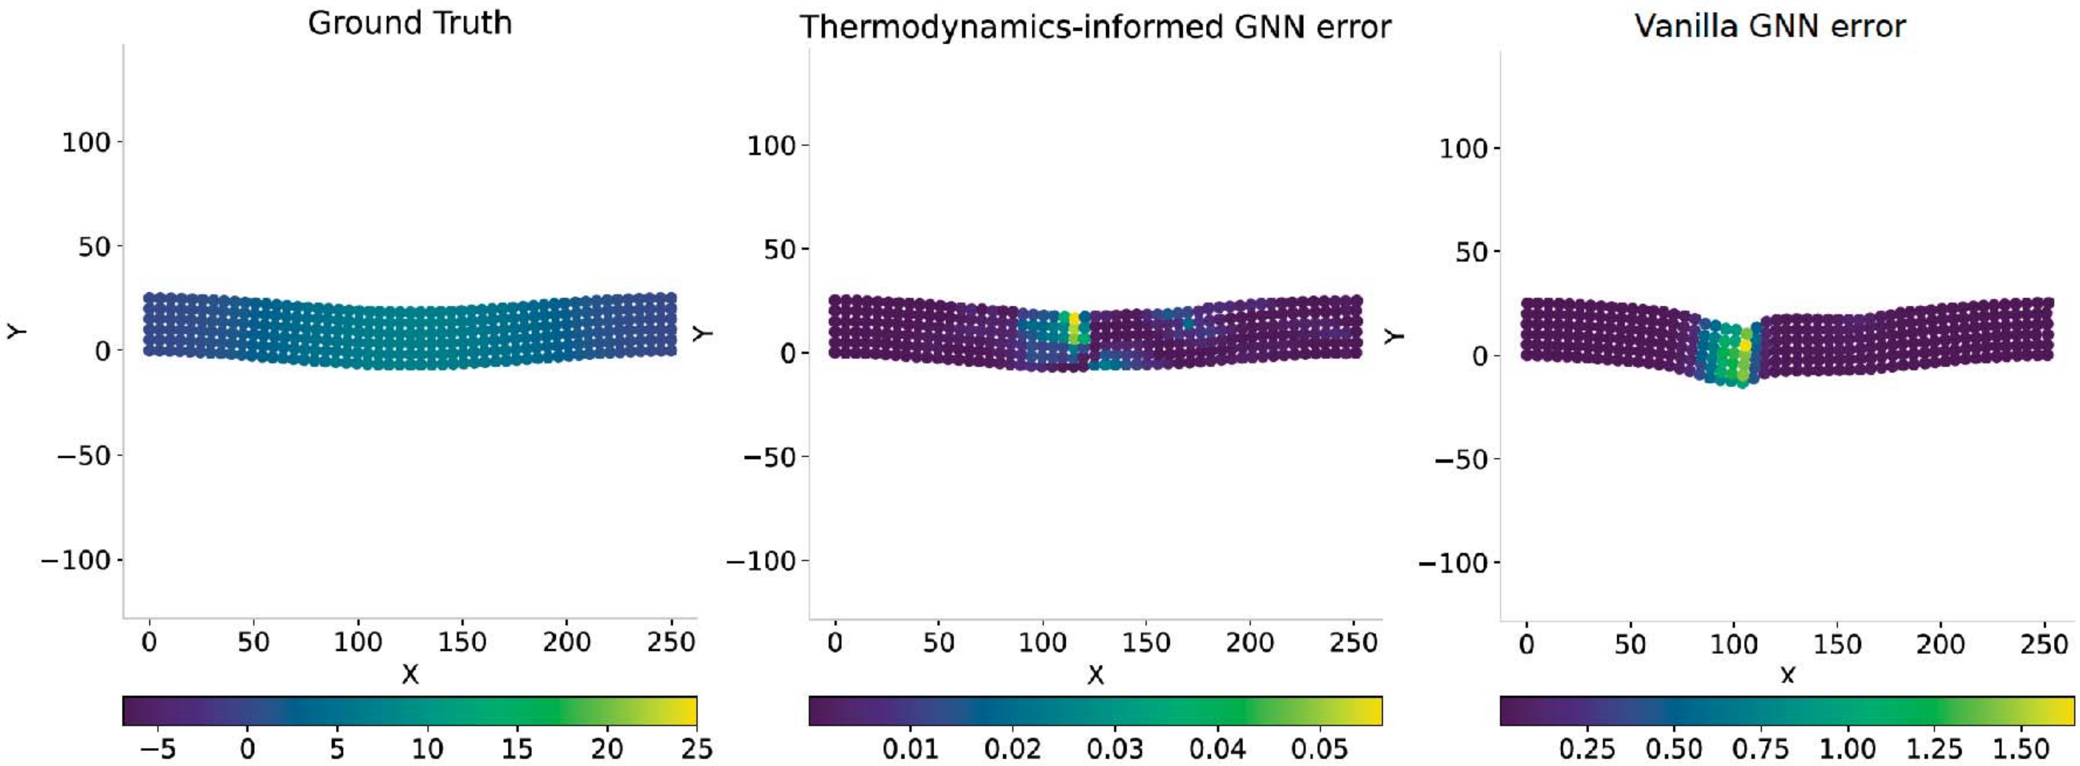

Fluids: previously unseen container geometry

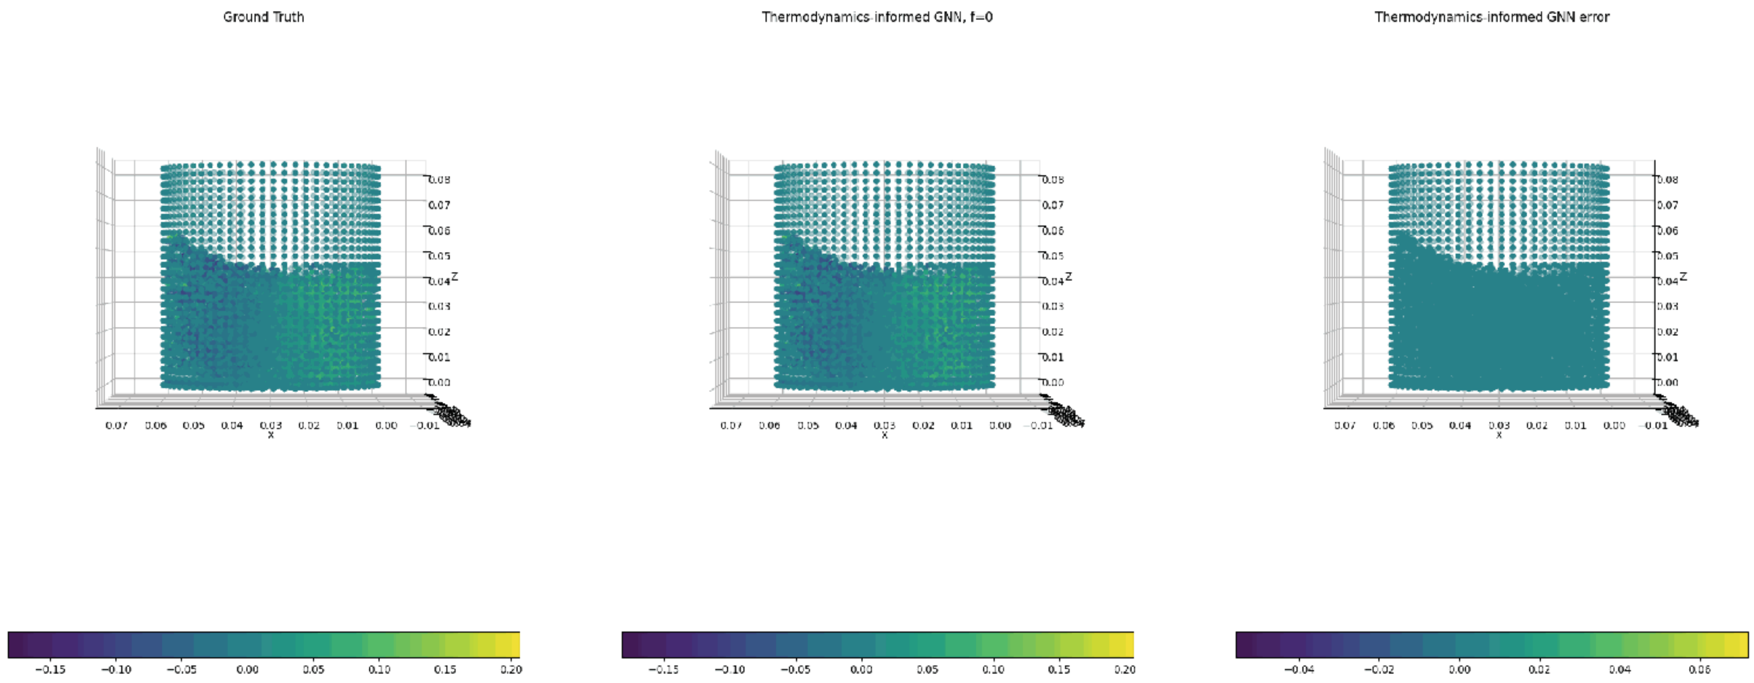

Tierz, Alicia, et al. "Graph neural networks informed locally by thermodynamics." *arXiv preprint arXiv:2405.13093* (2024).

Digital human twins

- Set of 20 actual liver anatomies provided by IRCAD, France

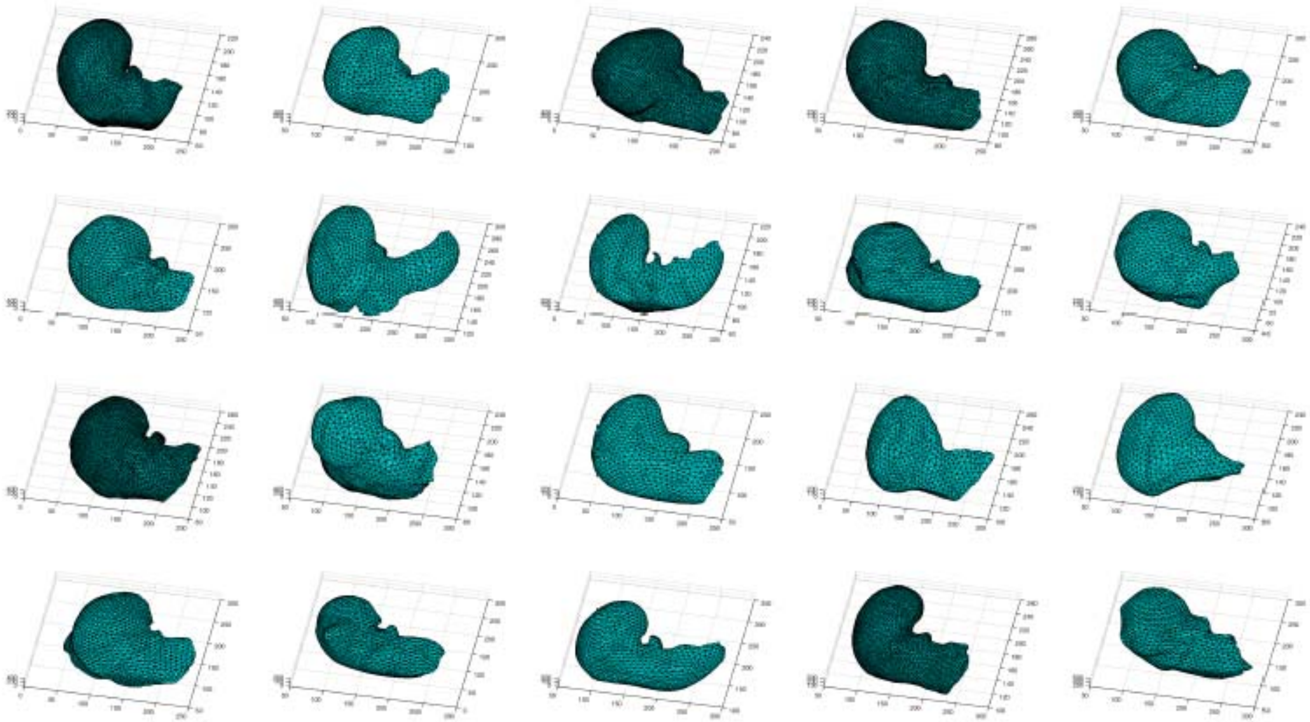

Ground truth vs. prediction

- Previously unseen anatomies

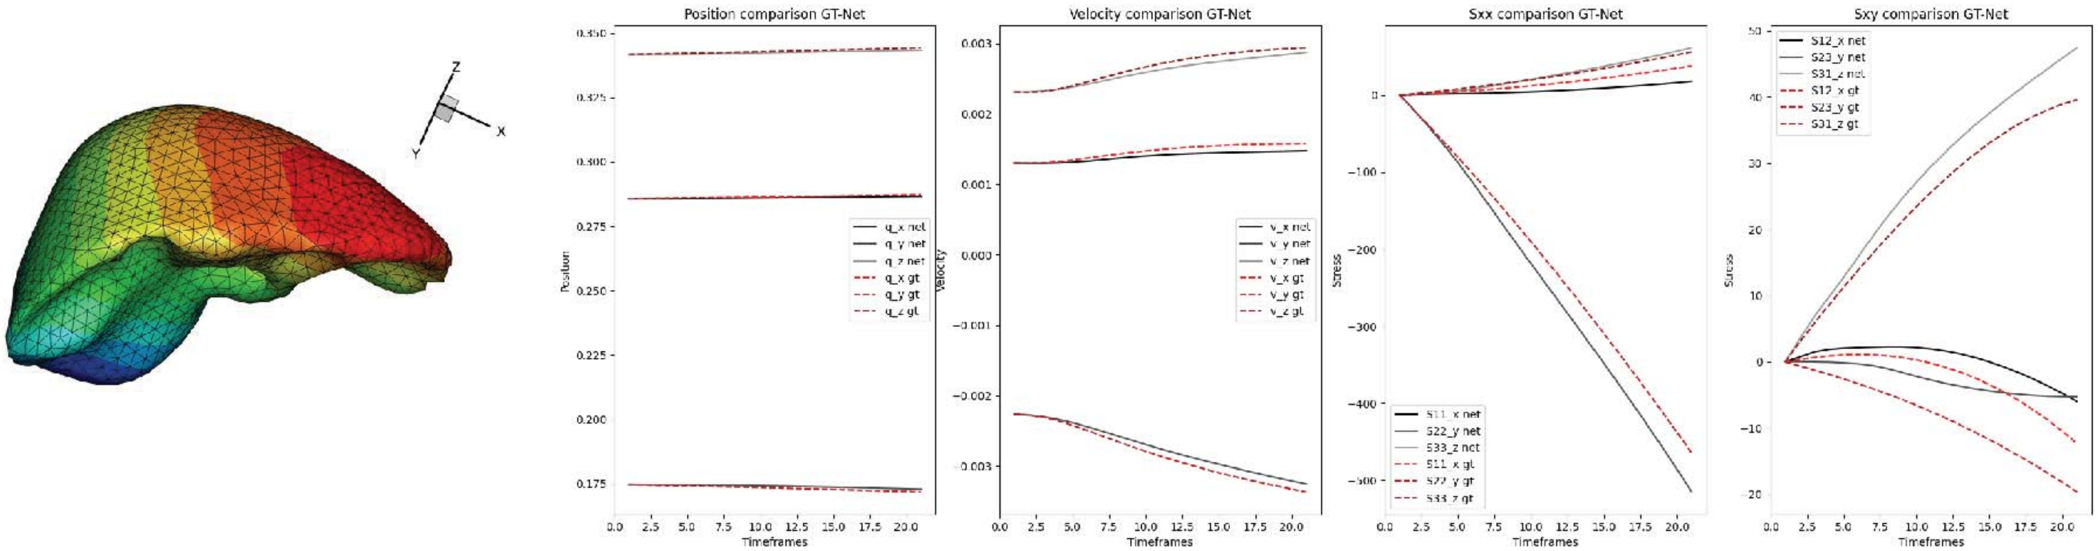

Digital human twins

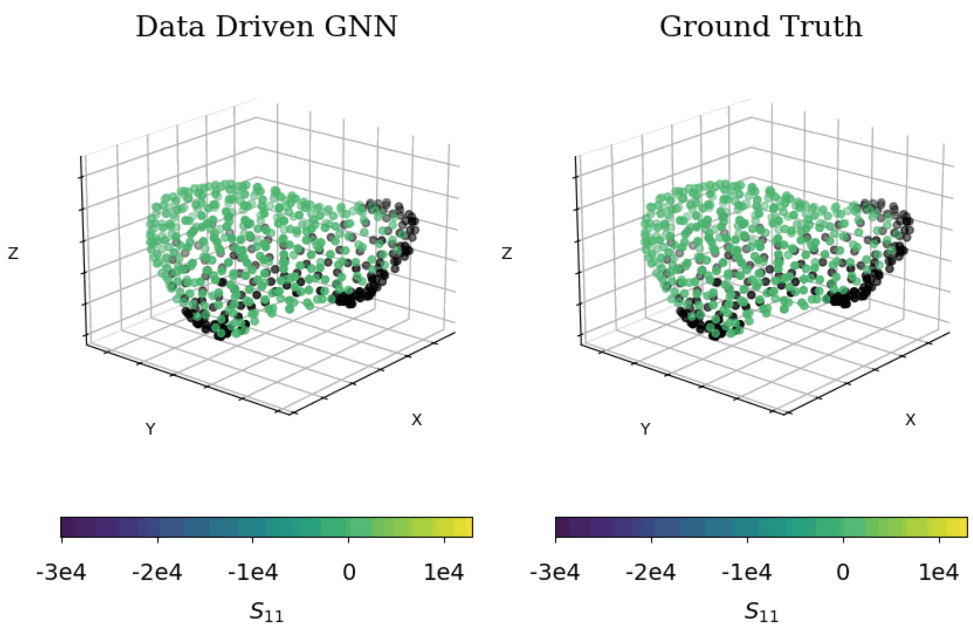

$$\|\boldsymbol{u}^{\mathrm{GT}} - \boldsymbol{u}\|_2 = 2.37 \cdot 10^{-3}.$$

Contents

Changes in constitutive equations

- We train our model with a Newtonian model (glycerine)
- Then, we face it against a (possibly non-Newtonian) different fluid
- Employ RL to let the system learn from observation
- Partial data regime: the camera only sees the free surface!

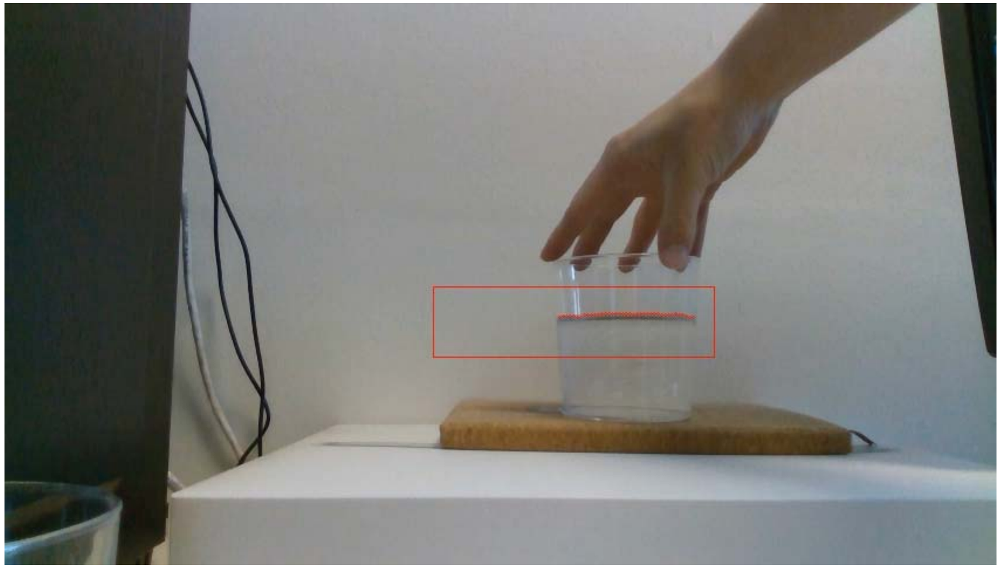

Perception + reasoning

Moya, Beatriz, et al. "Physics perception in sloshing scenes with guaranteed thermodynamic consistency." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 45.2 (2022): 2136-2150.

Thermodynamics-informed continuous learning

- What happens when the system is faced to previously unseen liquids?
- We employ reinforcement learning

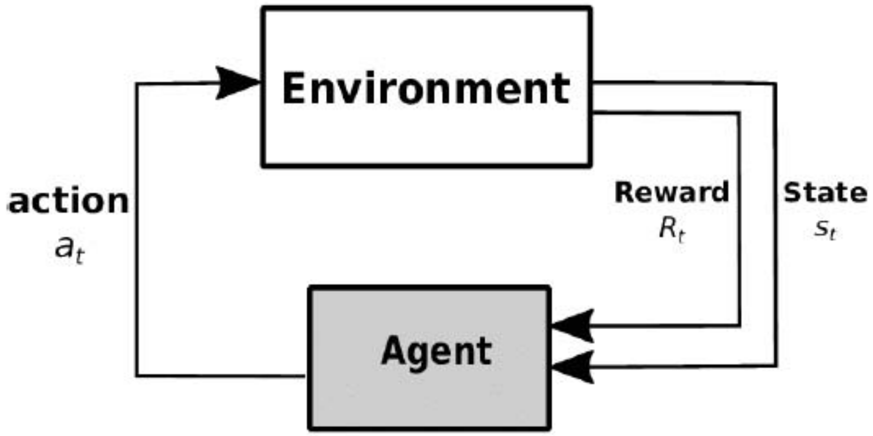

- Agent: our SPNN.
- Action: computation of GENERIC parameters.

Thermodynamics-informed continuous learning

- Reward

$$r_\pi = \lambda \frac{1}{N} \sum_N \left\| \boldsymbol{z}_{n+1} - \hat{\boldsymbol{z}}_{n+1} \right\|^2 + \frac{1}{N} \sum_N \left(\left\| \boldsymbol{L}_n \frac{\partial E_n}{\partial \boldsymbol{s}_n} \right\|^2 + \left\| \boldsymbol{M}_n \frac{\partial S_n}{\partial \boldsymbol{s}_n} \right\|^2 \right).$$

Thermodynamics-informed continuous learning

- Training with glycerine, then faced to water:

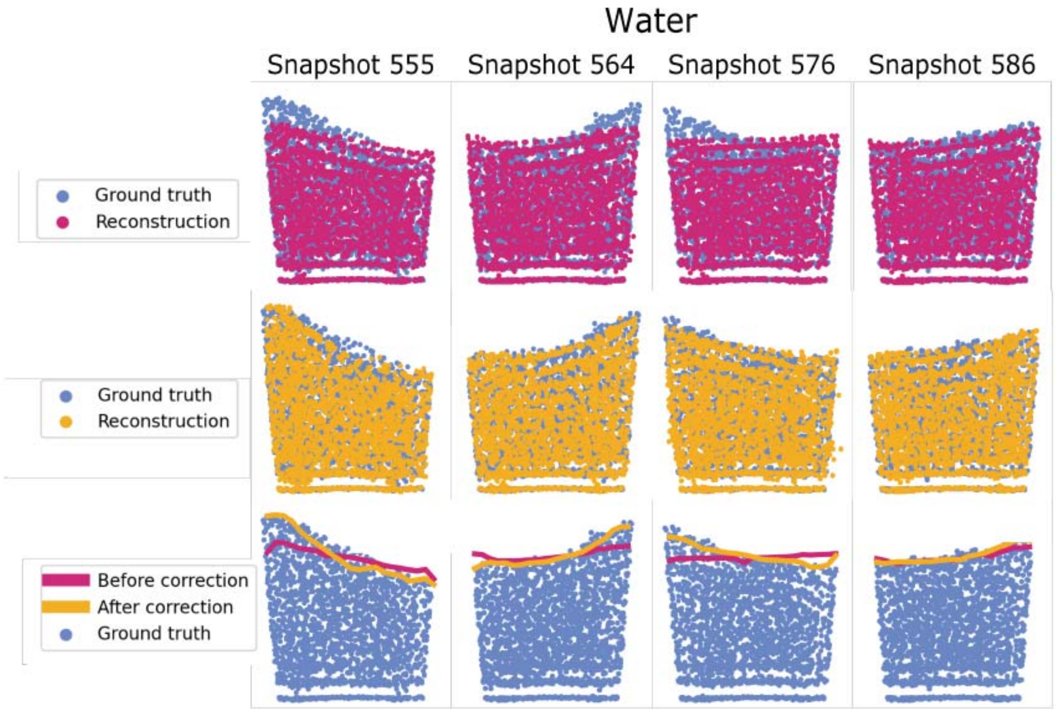

Moya, Beatriz, et al. "Computational Sensing, Understanding, and Reasoning: An Artificial Intelligence Approach to Physics-Informed World Modeling." *Archives of Computational Methods in Engineering* 31.4 (2024): 1897-1914.

Thermodynamics-informed continuous learning

- Training with glycerine, then faced to honey:

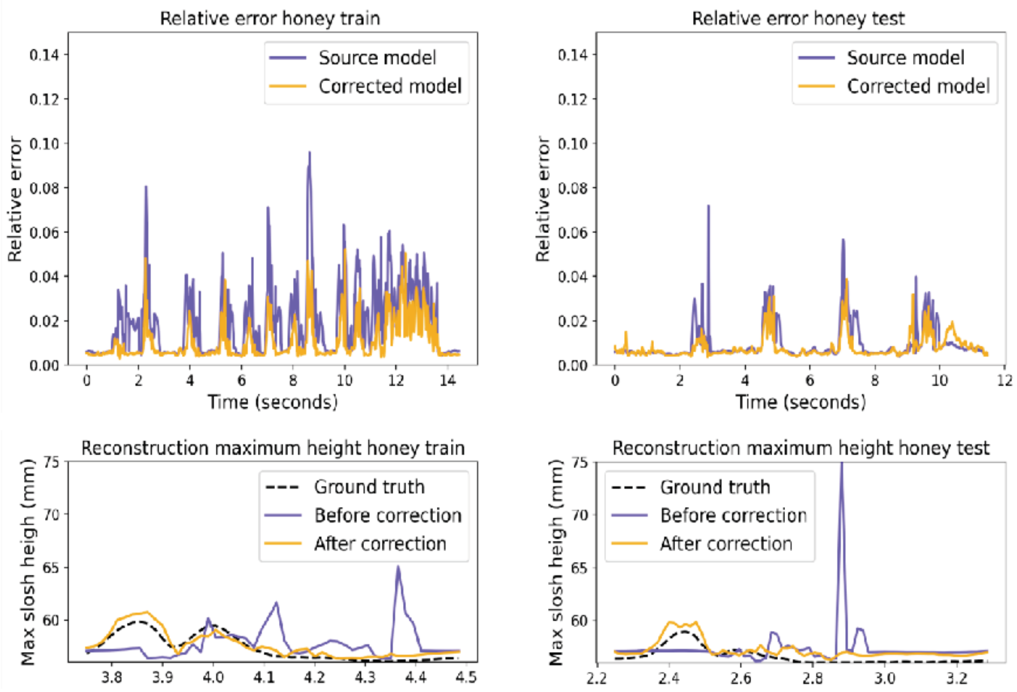

Moya, Beatriz, et al. "Computational Sensing, Understanding, and Reasoning: An Artificial Intelligence Approach to Physics-Informed World Modeling." *Archives of Computational Methods in Engineering* 31.4 (2024): 1897-1914.

Perception & reasoning

Reconstruction of state variables

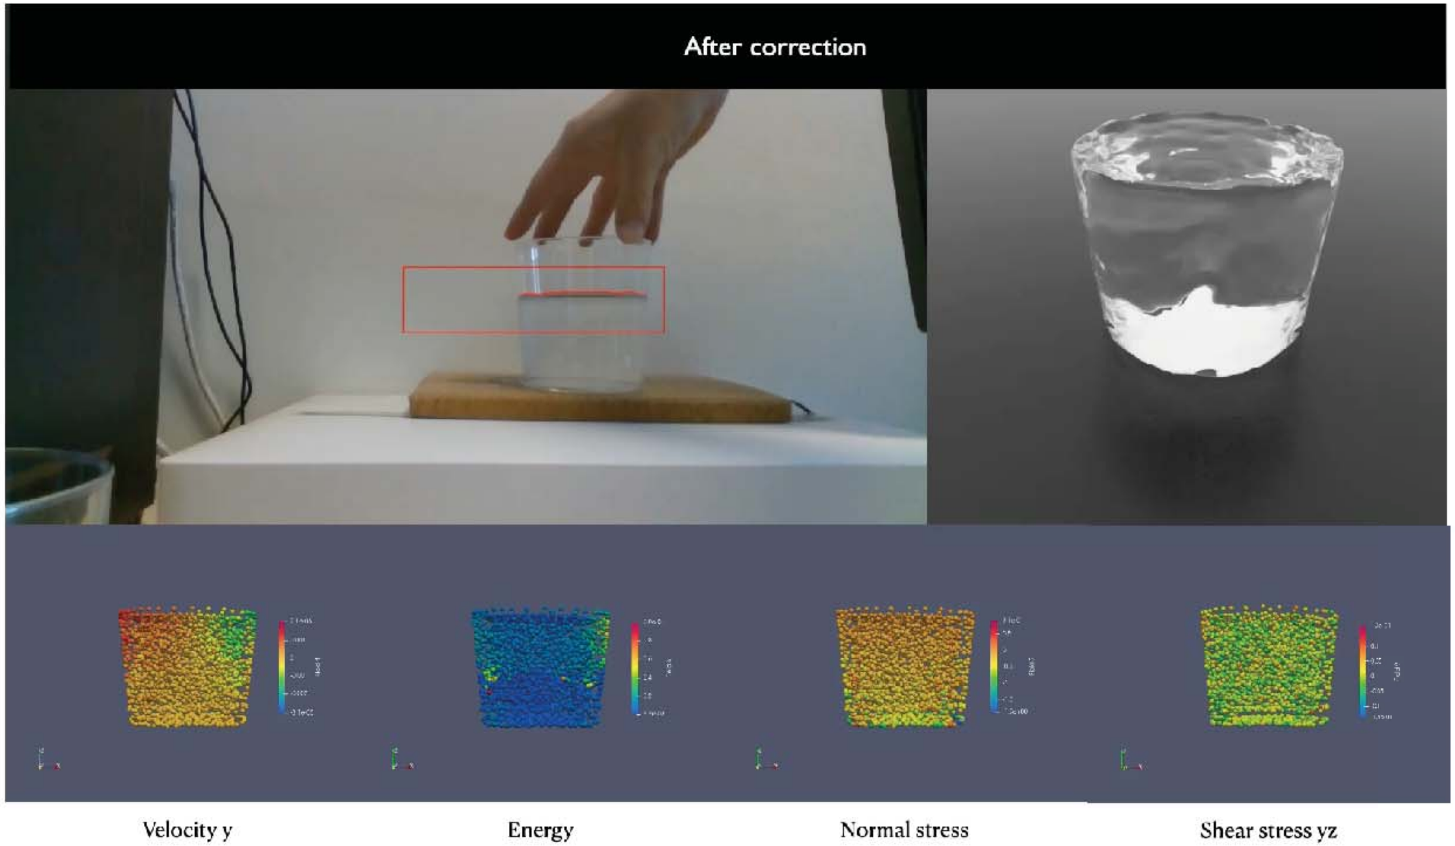

Contents

Conclusions

- Thermodynamics as inductive bias
- Robustness, accuracy in o.o.d. testing
- Thermodynamics-informed GNNs as a promising choice
- Size matters!

+info, preprints, …

https://eniachair.unizar.es

Acknowledgements

- The authors acknowledge the support of the Ministry for Digital Transformation and the Civil Service, through the ENIA 2022 Chairs for the creation of university-industry chairs in AI, through Grant TSI-100930-2023-1.
- The authors acknowledge the support by ESI Group through the ENIA Chair
- This material is based upon work supported in part by the Army Research Laboratory and the Army Research Office under contract/grant number W911NF2210271.