Physics-Informed Machine Learning Part III: Hard-Constraint ODE Method for Structural Dynamics

Mohsen Gol Zardian[1], Austin R.J. Downey[1,2], Eleonora Maria Tronci[3], Conor Madden[1], Daniel Coble[1,4], Sina Navidi[5], and Chao Hu[5]

[1]Department of Mechanical Engineering, University of South Carolina
[2]Department of Civil and Environmental Engineering
University of South Carolina
[3]Department of Civil and Environmental Engineering New York University
[4]Department of Mechanical Engineering & Materials Science Duke University
[5]School of Mechanical, Aerospace, and Manufacturing Engineering University of Connecticut

UNIVERSITY OF SOUTH CAROLINA

Physics-Informed Machine Learning (PIML): Big Picture

- Combines governing physics with data-driven learning
- Reduces data requirements compared to data-driven models
- Improves physical consistency and interpretability

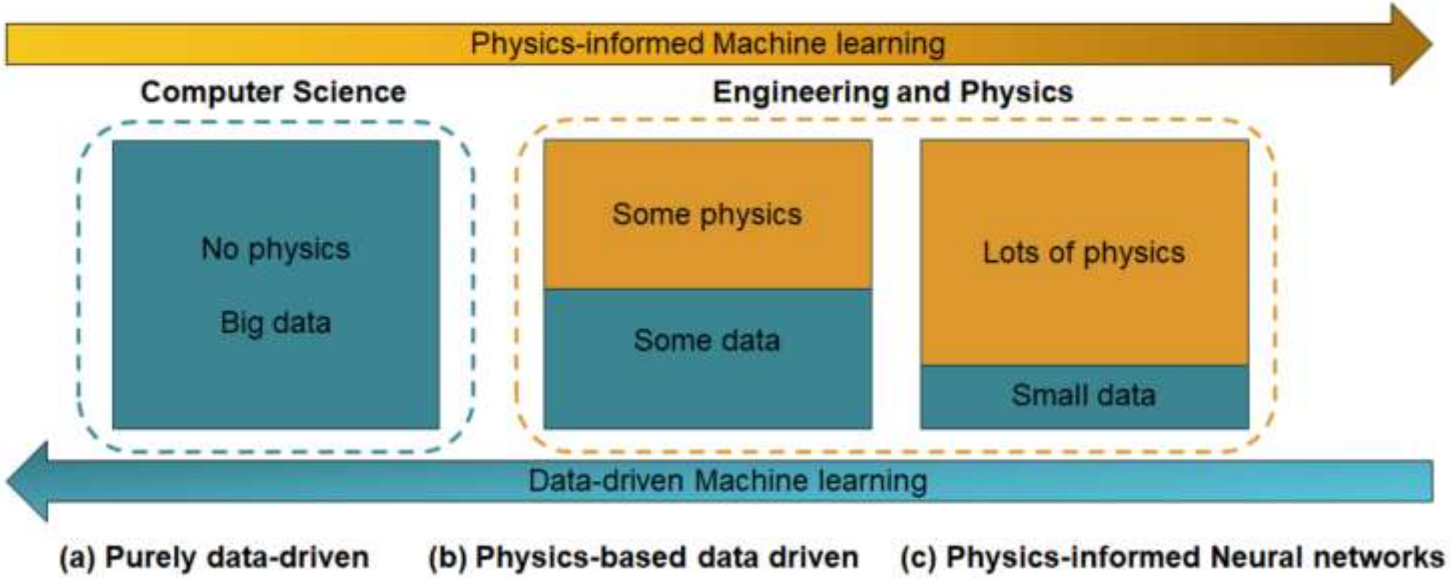

Puja Chowdhury. Temporal Forecasting of High-rate Dynamics using Physics-informed Machine Learning and Hardware-software Co-design. University of South Carolina, 2024, University of South Carolina Graduate Theses and Dissertations.

Why Purely Data-Driven Models Fall Short

- Require large, high-quality datasets
- Can violate known physical laws
- Poor extrapolation beyond training regime
- Limited interpretability for engineering decisions

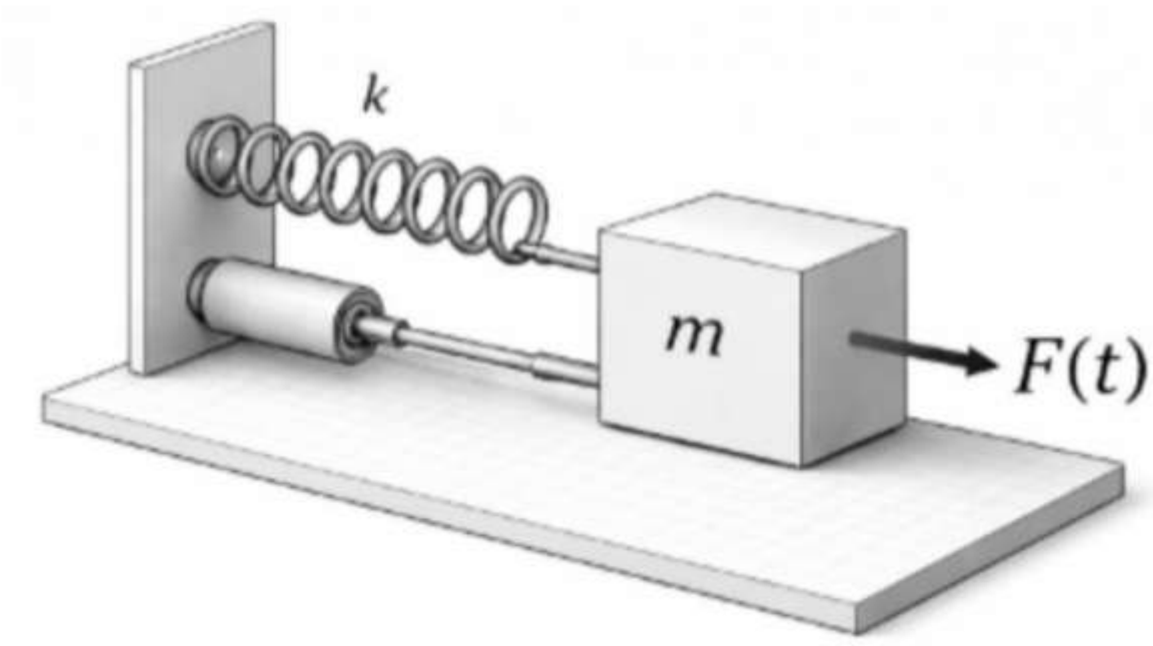

What PIML Adds to the Learning Process

- Encodes governing equations into learning
- Constrains predictions to physically admissible solutions
- Improves generalization under noise and sparsity
- Bridges modeling and sensing

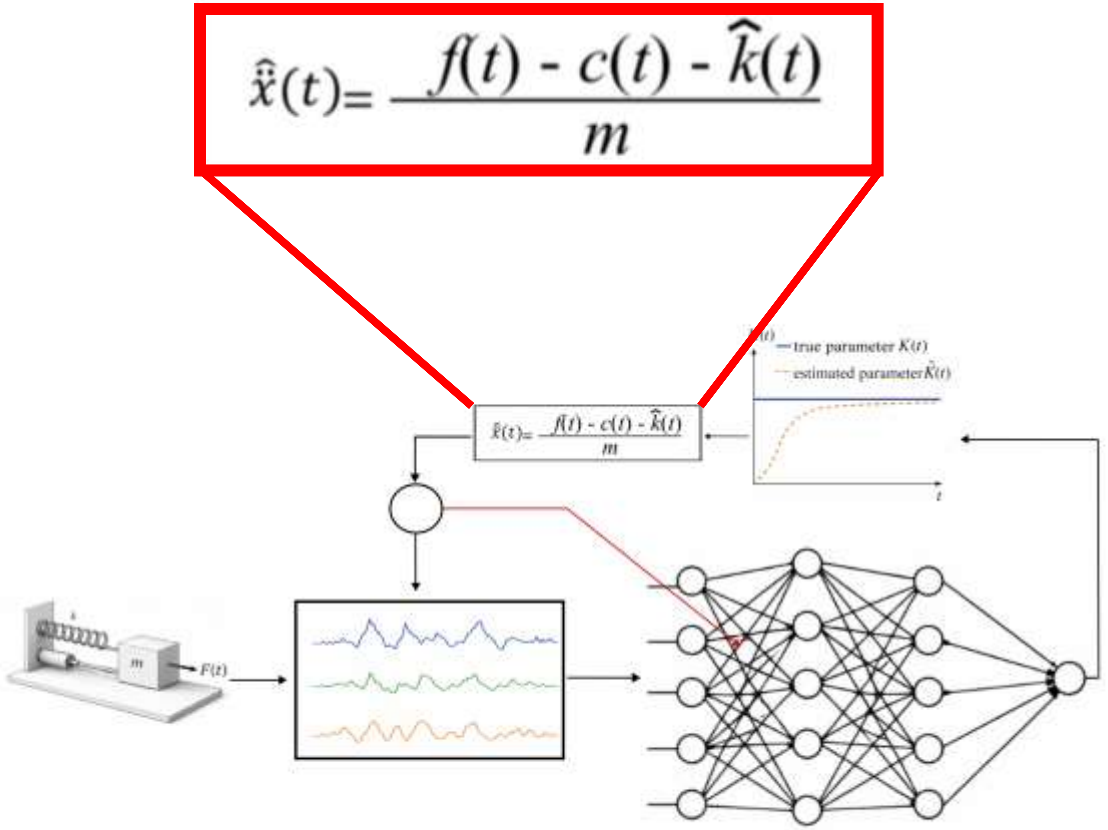

Eleonora Maria Tronci, Austin R.J. Downey, Azin Mehrjoo, Puja Chowdhury, and Daniel Coble. Physics informed machine learning part I: Different strategies to incorporate physics into engineering problems. In Conference Proceedings of the Society for Experimental Mechanics Series. Springer Nature Switzerland, 2024. doi:10.1007/978-3-031-68142-4_1

Two Classes of PIML Approaches

- Hard-constraint PIML (Part III)
  - Physics embedded directly in model architecture
- Soft-constraint PIML (Part IV)
  - Physics residual added to loss function
- Trade-off between flexibility and strict enforcement

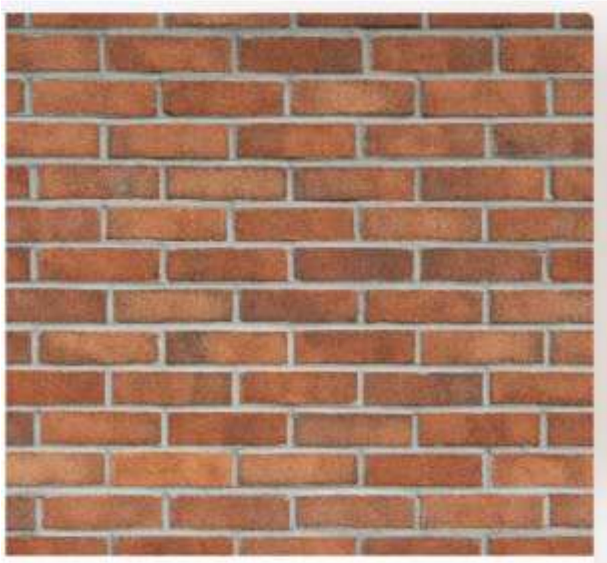

Can’t go through, no matter the penalty

Go through if willing to pay the penalty

Part II Recap: Structural Response Forecasting

- Combines governing physics with data-driven learning
- Reduces data requirements compared to purely data-driven models
- Improves physical consistency and interpretability
- Especially valuable for dynamical systems

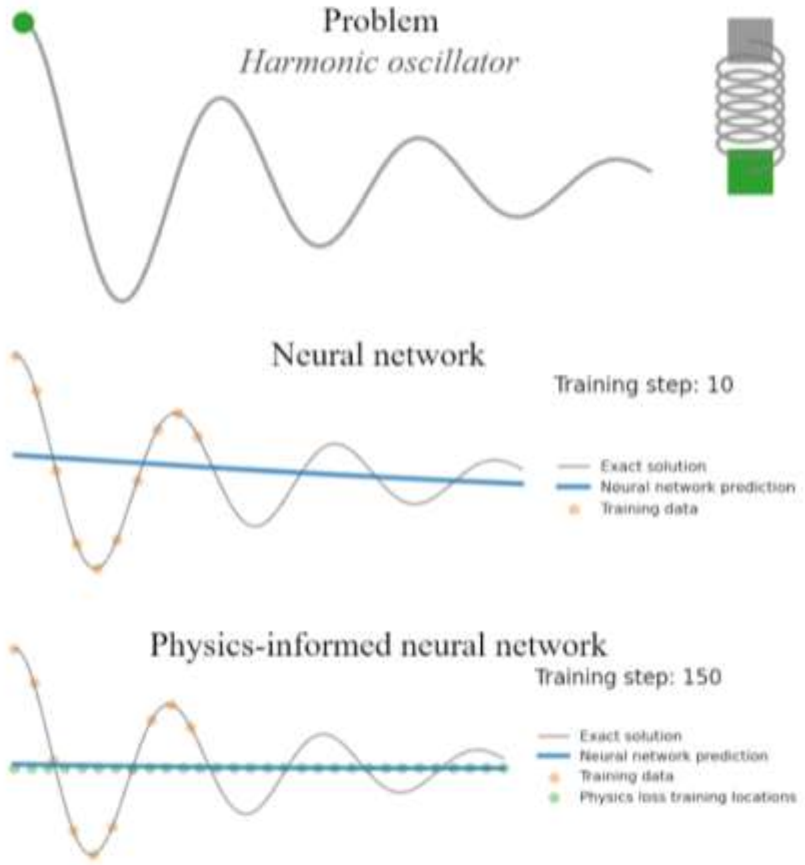

Austin R.J. Downey, Eleonora Maria Tronci, Puja Chowdhury, and Daniel Coble. Physics informed machine learning part II: Applications in structural response forecasting. In *Conference Proceedings of the Society for Experimental Mechanics Series*. Springer Nature Switzerland, 2024. doi:10.1007/978-3-031-68142-4_8

https://www.reddit.com/r/learnmachinelearning

Part III: Shift from Forecasting to State Estimation

- Focus shifts from prediction to state estimation
- Emphasis on inverse problem
- Estimate hidden system properties
- Physics must be satisfied at every step

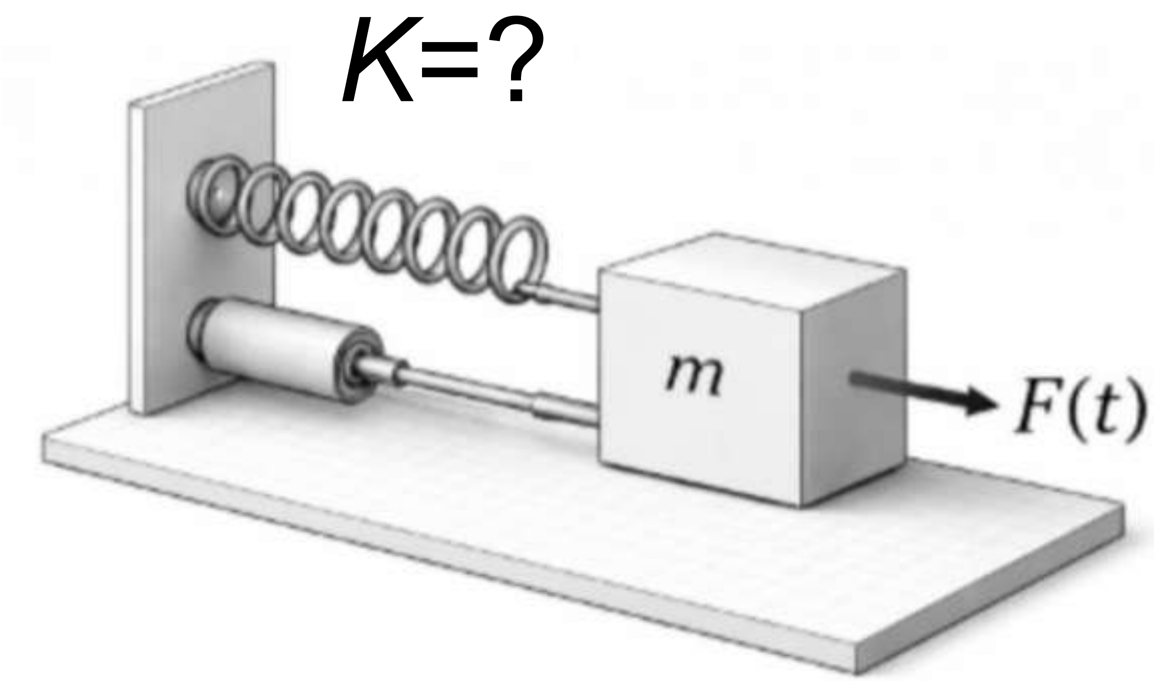

Targeted Study

Targeted Study Overview

- Demonstration of hard-constraint PIML on a canonical system
- Focus on inverse problem: stiffness estimation
- Controlled numerical study to isolate behavior
- Benchmark against physics-only and data-only models

Physics-Informed Machine Learning Part III: Hard-Constraint ODE Method for Structural Dynamics

Mohsen Gol Zardian[1], Austin R.J. Downey[1,2], Eleonora Maria Tronci[3], Conor Madden[1], Daniel Coble[1,4], Sina Navidi[5], and Chao Hu[5]

[1]Department of Mechanical Engineering
University of South Carolina, Columbia, SC 29208

[2]Department of Civil and Environmental Engineering
University of South Carolina, Columbia, SC 29208

[3]Department of Civil and Environmental Engineering
New York University
[4]Department of Mechanical Engineering & Materials Science
Duke University
[5]School of Mechanical, Aerospace, and Manufacturing Engineering
University of Connecticut

Abstract

Physics-informed machine learning (PIML) is a methodology that combines principles from physics with machine learning (ML) techniques to enhance the accuracy and interpretability of predictive models. By incorporating physical laws and constraints into the learning process, physics-informed machine learning enables more robust predictions and reduces the need for large amounts of training data. PIML has a wide range of applications in science and engineering, such as modeling physical systems, solving partial differential equations, and performing inverse analysis and optimization.

In Part III of this series, the authors provide an overview of a hard-constraint, indirect-measurement PIML approach for structural dynamics. An ordinary differential equation (ODE) solver implemented within the automatic differentiation engine serves as a differentiable layer that enforces the mass-damping-stiffness equilibrium at every inference step. A compact neural network ingests a window of past acceleration samples and outputs an effective stiffness. The embedded Runge-Kutta-4 integrator then propagates states so that the predicted acceleration satisfies the equation of motion. The technique is demonstrated on a deliberately simple numerical example: a simulated spring, mass, and dashpot whose stiffness drifts over time along several representative paths. A collection of runs is generated, some used for training, the rest held back for validation so that the model must generalize to unseen stiffness trajectories. Performance is benchmarked against two reference solvers. First, a physics-only baseline consisting of a conventional ODE integrator that assumes the stiffness is fixed and therefore cannot reflect its time variation. Second, a data-only baseline consisting of a neural network of identical size trained to minimize next-step acceleration error, without any governing-equation guidance. Results show that the hard-constraint PIML model approaches the data-only network in response accuracy while offering far greater physical interpretability through its explicit stiffness estimates. The paper reports on hyper-parameter tuning for the multi-layer perceptron so readers can readily adapt the hard-constraint indirect-measurement PIML method to their own inverse-problem applications. Project code are artifacts available are through thorough a public repository.

Keywords: Physics-informed, physics-based model, data driven model, machine learning, stiffness identification

Benchmark Structural System

- Single-degree-of-freedom spring–mass–damper system
- Widely used model in structural dynamics
- Simple enough for interpretability
- Rich enough to expose inverse-problem challenges

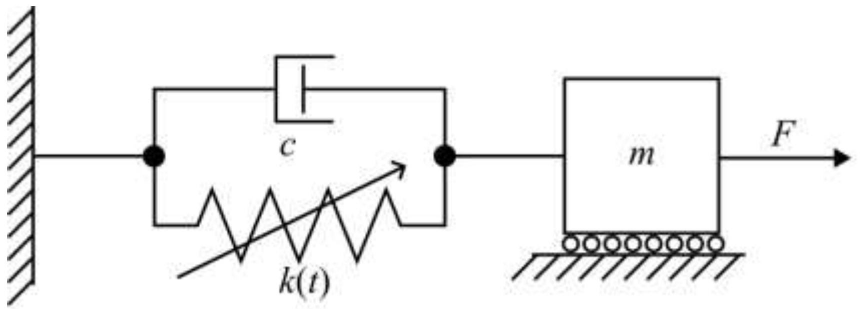

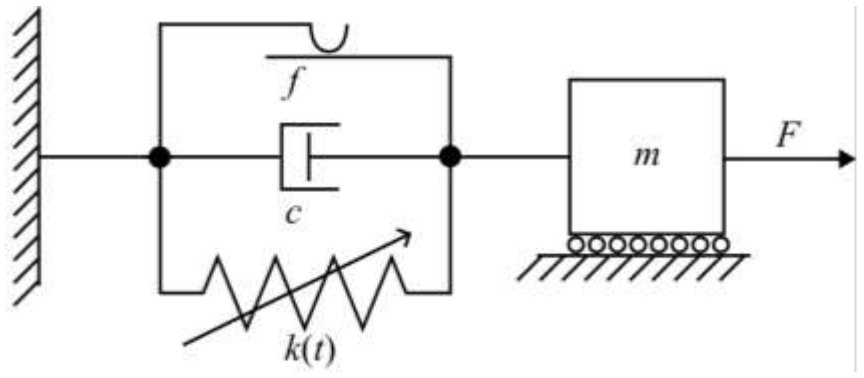

$$m\ddot{x} + c\dot{x} + kx = 0$$

Time-Varying Stiffness and Resonance Behavior

- Prescribed stiffness decreases over time
- Reduction in stiffness lowers natural frequency
- System passes through resonance as frequencies align
- Produces strong, nonstationary response features

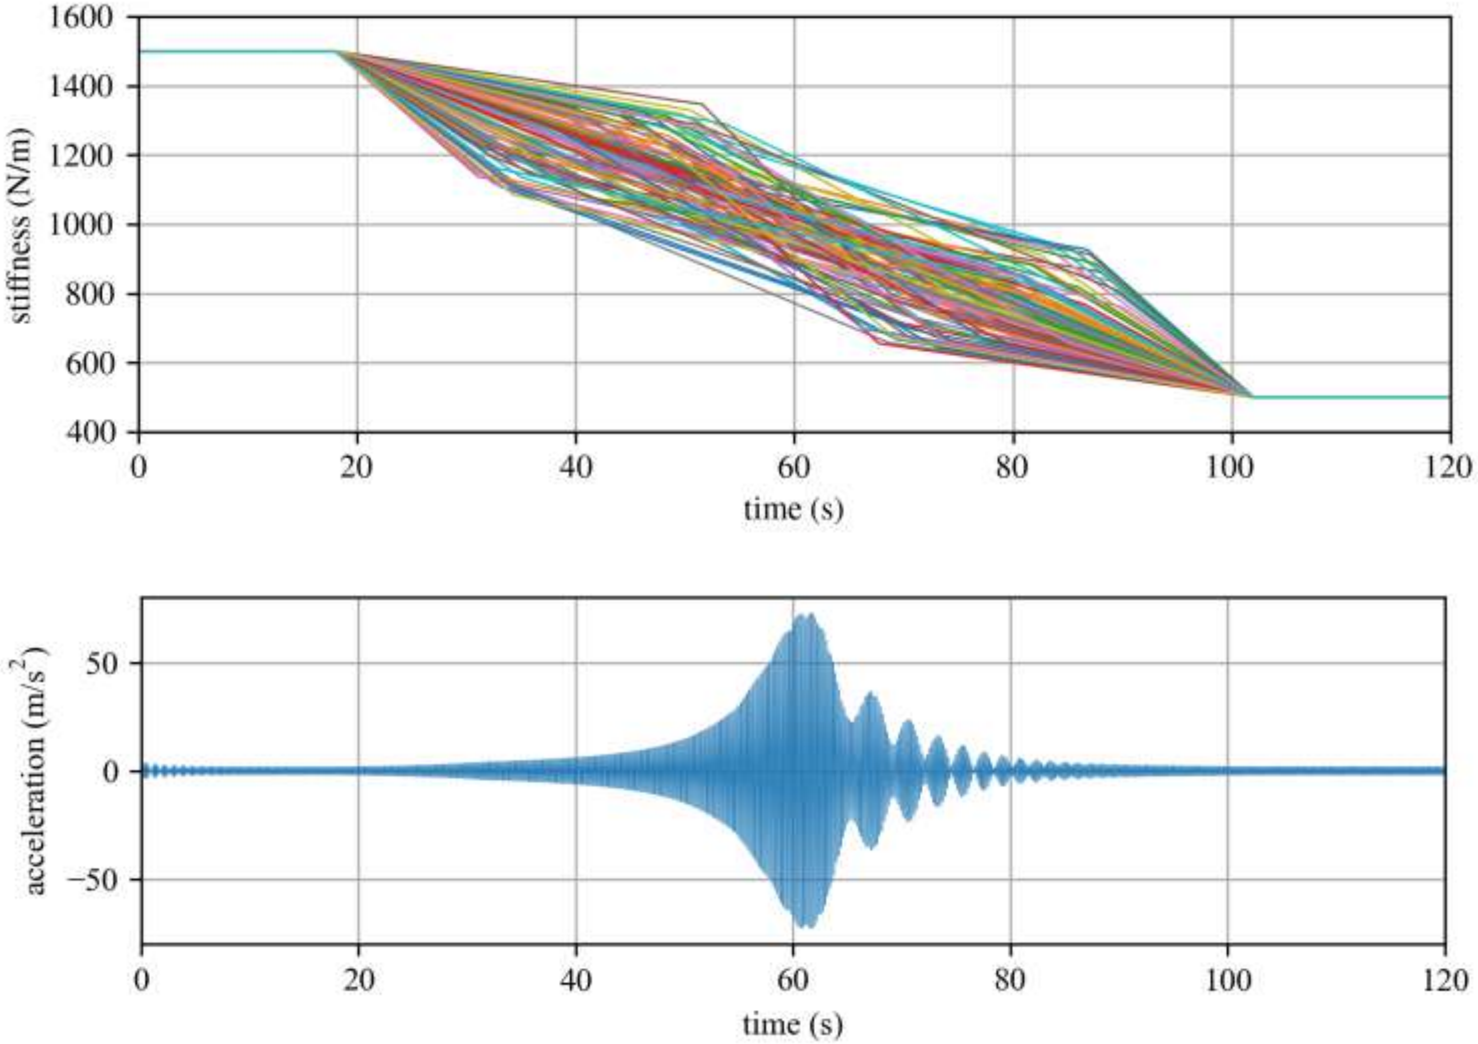

Why This Becomes an Inverse Problem

- Stiffness is not directly measured
- Only acceleration response is observed
- Multiple stiffness histories can explain similar responses

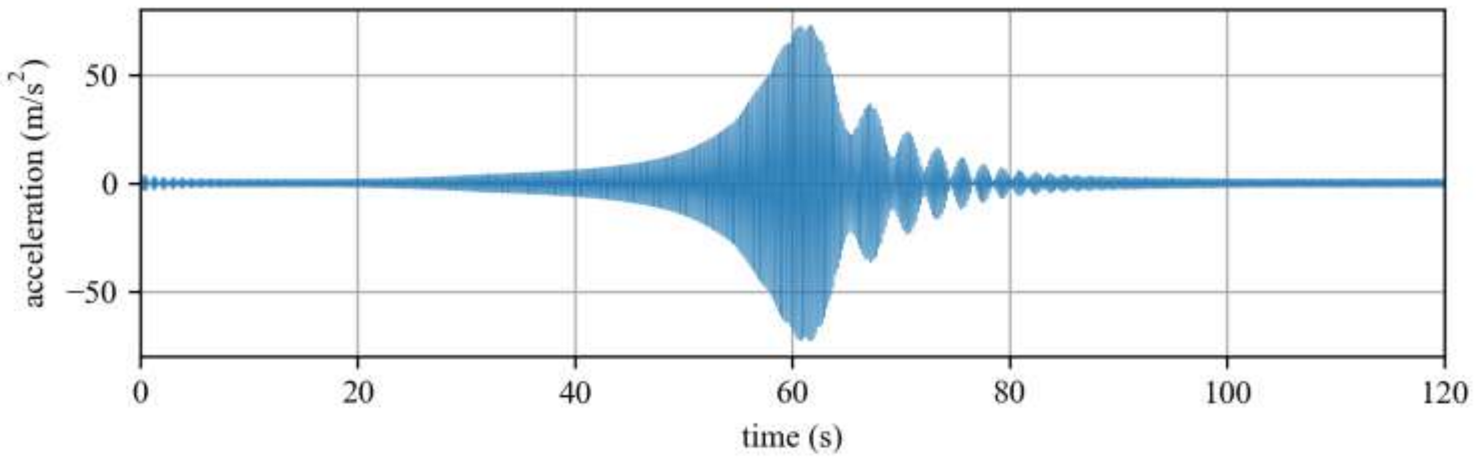

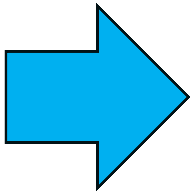

K=?

Hard-Constraint PIML: Key Idea

- Neural network estimates stiffness
- Governing ODE checks k by solving the ODE at every inference step
- No physics residual added to the loss (just RMSE)
- Physics enforced in model architecture

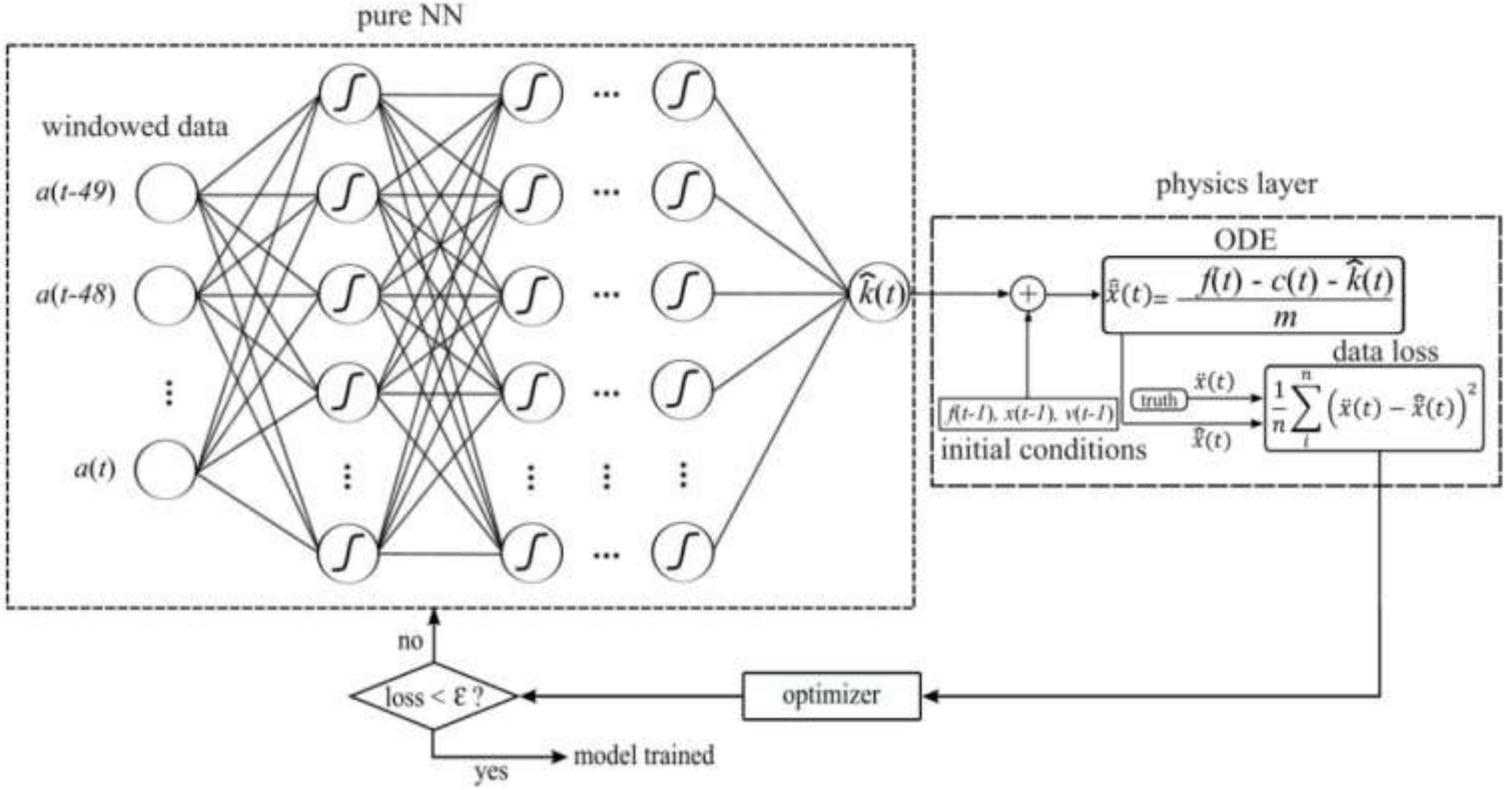

$$\mathscr{L} = \frac{1}{N}\sum_{i=1}^{N}\left(\ddot{x}(t)_i - \hat{\ddot{x}}(t)_i\right)^2$$

What the Model Sees

- Input: windowed acceleration history
- No access to displacement or stiffness
- Windowed acceleration fed into NN as input vector and the estimated stiffness is the output.
- Short time windows capture local dynamics

Algorithm 1 preprocessing pipeline from raw data to mini-batches

procedure BUILDBATCHES($\{\mathrm{CSV}_n\}_{n=1}^{N}, s, T, W, B, \rho$)

Require: each CSV_n has L rows and D columns (e.g., t, x, v, a, k, F)

Ensure: iterable of train/validation mini-batches with shape (B, T, W)

    counts/params: $N \to$ number of trajectories; $L \to$ rows/trajectory; $D \to$ columns; $T \to$ segment length; $W \to$ window length; $B \to$ batch size

    shape $N \times L \times D$**:** $X \leftarrow \mathrm{Stack}(\mathrm{CSV}_1, \dots, \mathrm{CSV}_N)$

    downsample by factor s**:** $\tilde{X} \leftarrow \mathrm{Downsample}(X, s)$; $\tilde{L} \leftarrow \mathrm{ceil}(L/s)$

    for $n = 1$ **to** N **do**

        for $t = W$ **to** $\tilde{L}$ **do**

            window length W**:** $a_{t:W}^{(n)} \leftarrow [a_{t-W+1}^{(n)}, \dots, a_t^{(n)}]$

        end for

        for $\tau = 1$ **to** $\tilde{L} - T + 1$ **do**

            segment size $T \times W$**:** $S_\tau^{(n)} \leftarrow (a_{\tau:W}^{(n)}, \dots, a_{\tau+T-1:W}^{(n)})$

        end for

    end for

    total segments: $N_{\mathrm{seg}} \leftarrow N(\tilde{L} - T + 1)$

    split ratio ρ**:** $\{\mathrm{train}\}, \{\mathrm{val}\} \leftarrow \mathrm{Split}(\{S_\tau^{(n)}\}, \rho)$

    tensor shape (B, T, W)**:** $\mathrm{batches} \leftarrow \mathrm{Batch}(\{\mathrm{train/val}\}, B)$

end procedure

Neural Network: Parameter Estimation

- Neural network ingests windowed acceleration history
- Outputs time-varying effective stiffness *k(t)*
- No prediction of acceleration
- Learns parameters, not dynamics

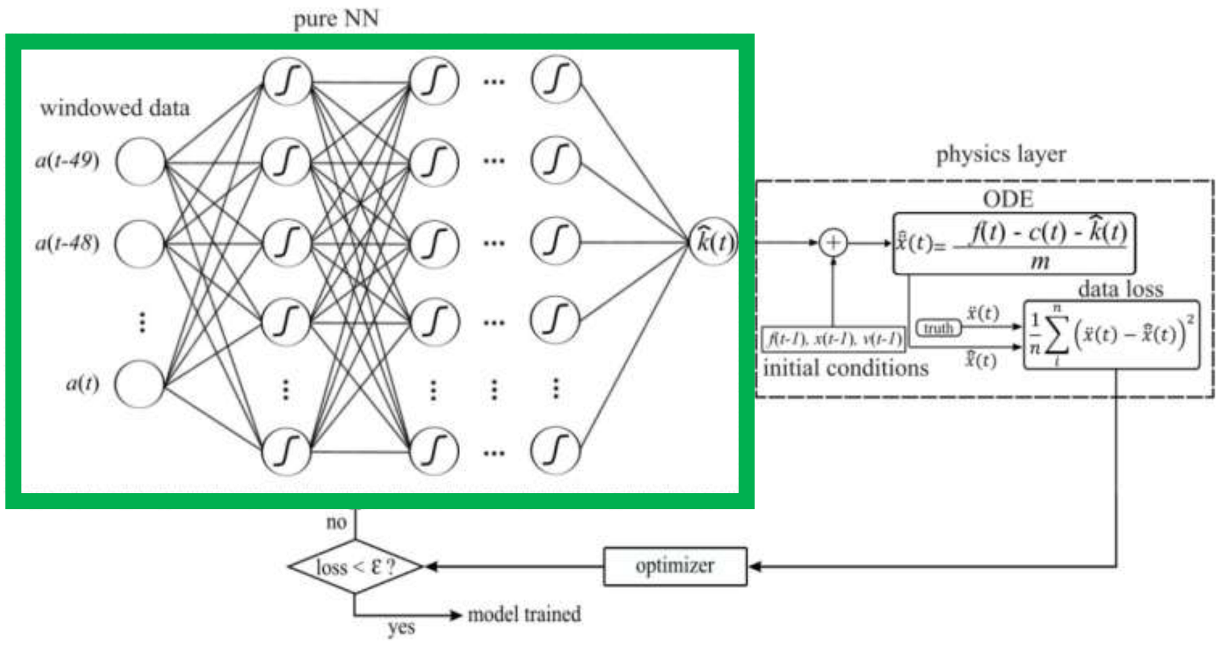

Key message: NN ≠ dynamics solver

Physics Layer: State Propagation

- Estimated stiffness passed to ODE solver
- Governing equation enforced exactly
- Displacement and velocity propagated in time
- Acceleration reconstructed from ODE states

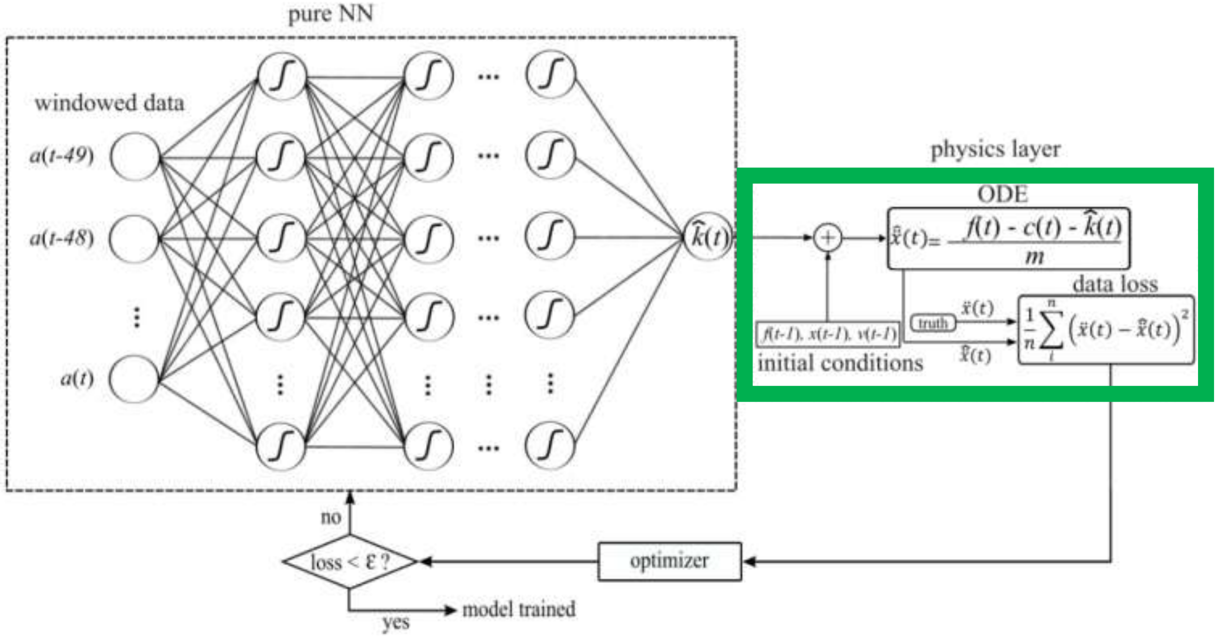

Key message: dynamics come from physics, not the NN

Training Loop: How Learning Happens

- Predicted acceleration compared to measured acceleration
- Loss defined as acceleration mismatch
- Error backpropagated through ODE solver
- Neural network updated to improve stiffness estimates

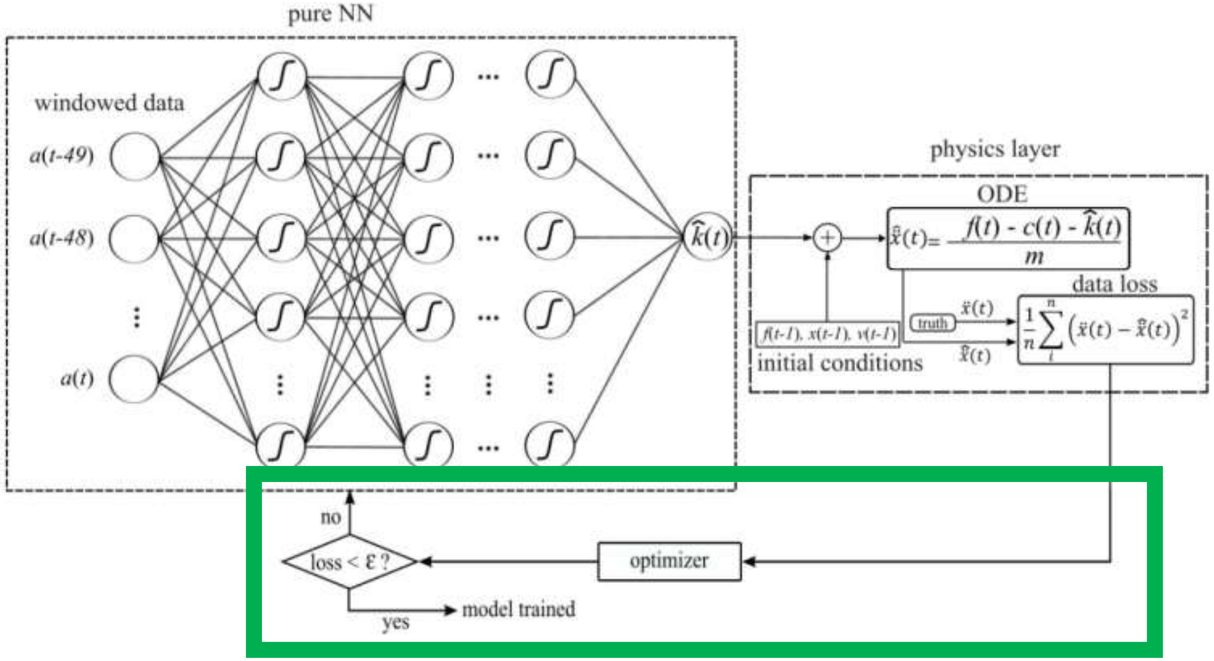

Key message: data trains the NN *through* physics

Inference: Hard-Constraint PIML Forward Pass

- NN estimates time-varying stiffness from acceleration windows
- ODE propagates displacement and velocity states forward in time
- Acceleration is reconstructed from the physics-based state update
- Governing dynamics are enforced at every inference step

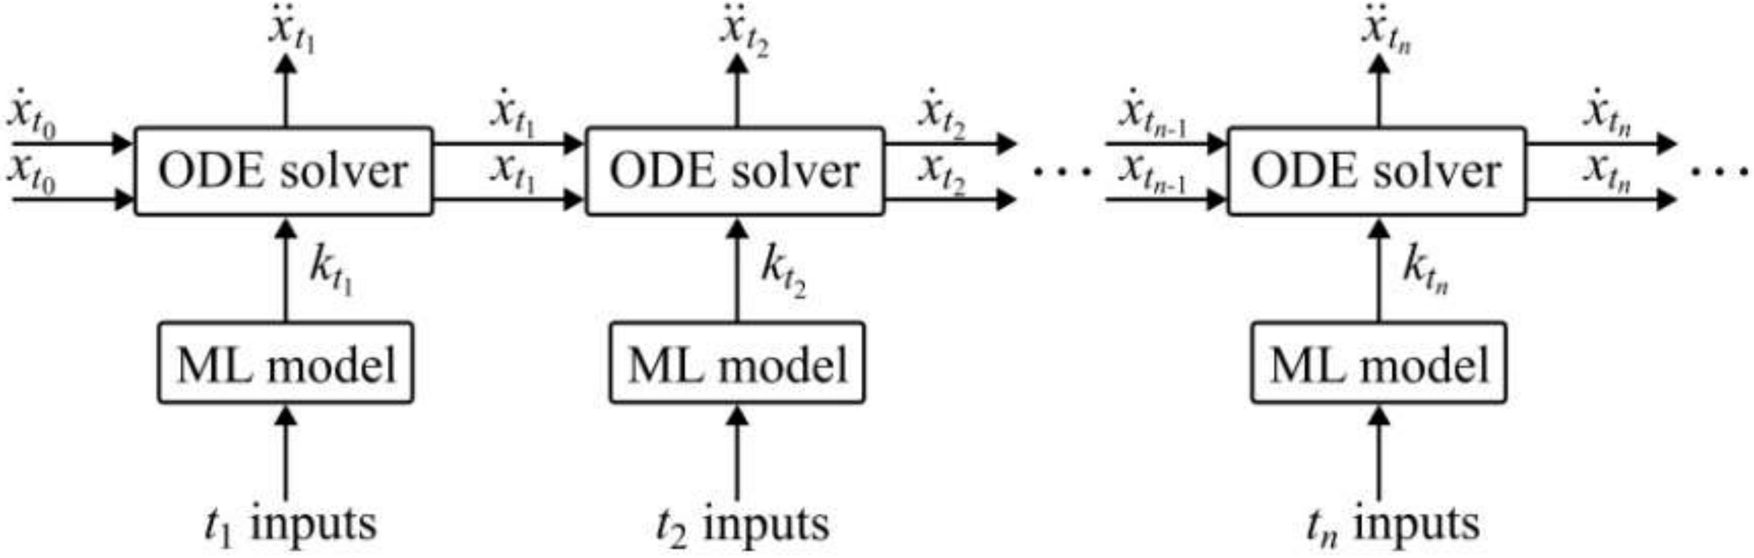

Results: Inverse Stiffness Identification (The Good)

- Time-varying stiffness inferred from acceleration-only measurements
- Hard-constraint PIML tracks global stiffness degradation
- Comparable accuracy to data-only model

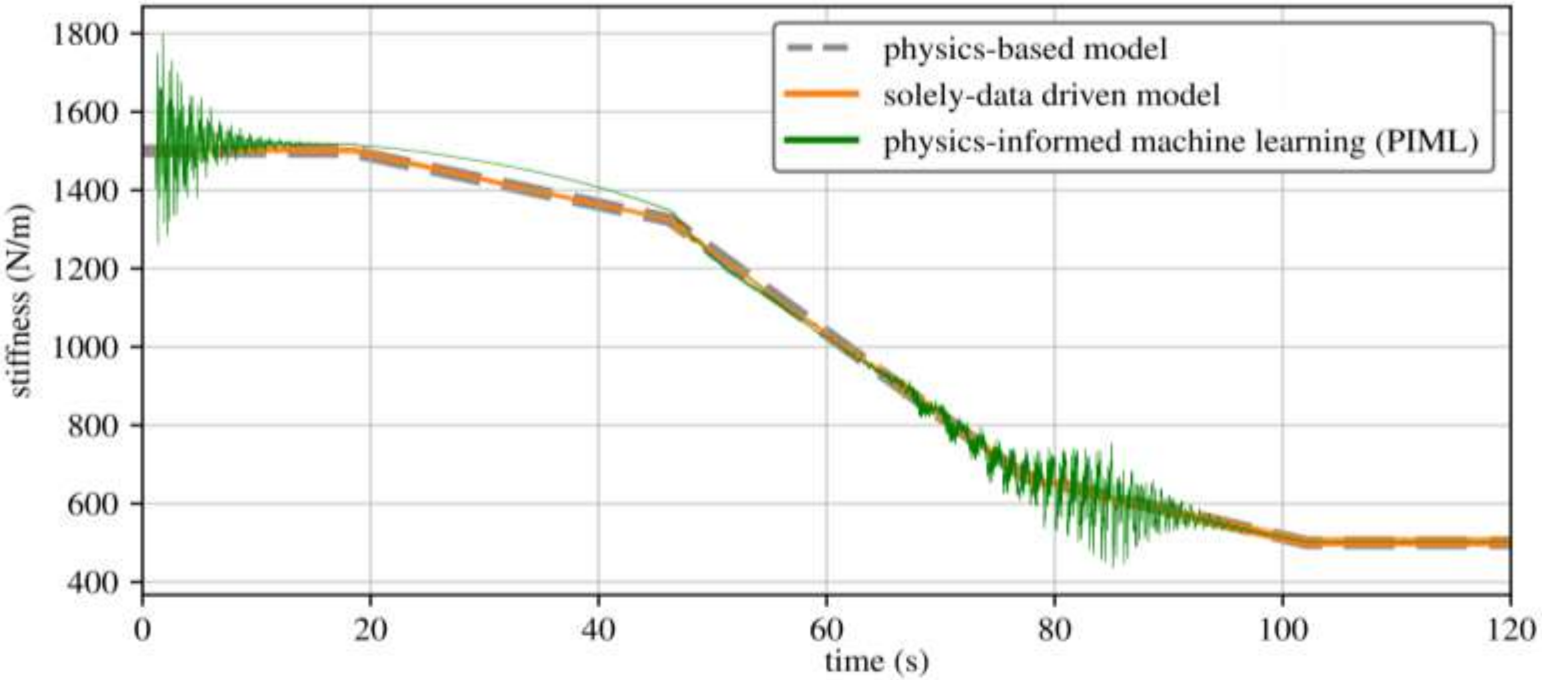

Results: Inverse Stiffness Identification (The Bad)

- High-frequency ripple reflects strict physics enforcement
- Computationally expensive due to having to solve the ODE at each step
- Performance depends on the correctness of the physical model

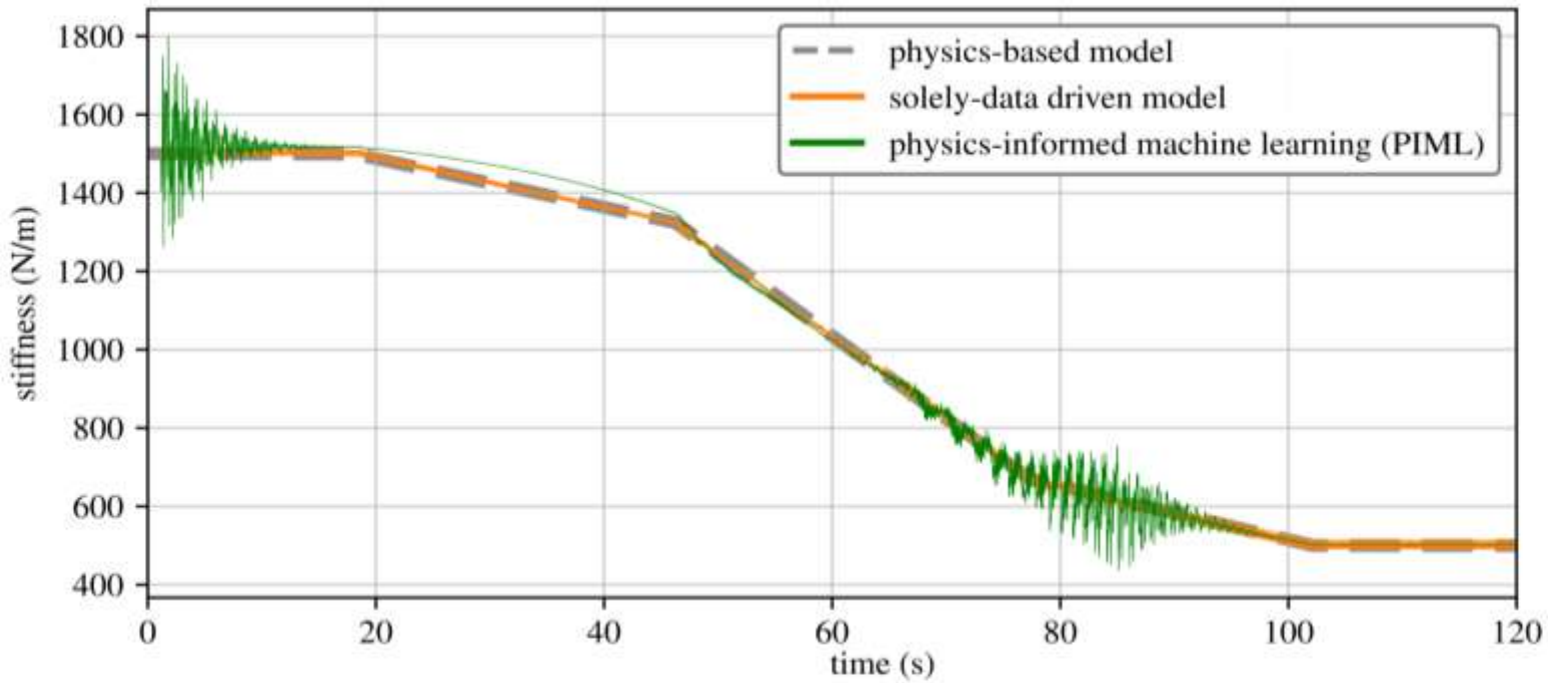

Activation Functions: Intuition

- Activation function controls how signals propagate through the network
- Sigmoid compresses outputs to [0,1]
- Tanh maps inputs to [−1,1 and is zero-centered
- Vibration signals are +/– and zero centered

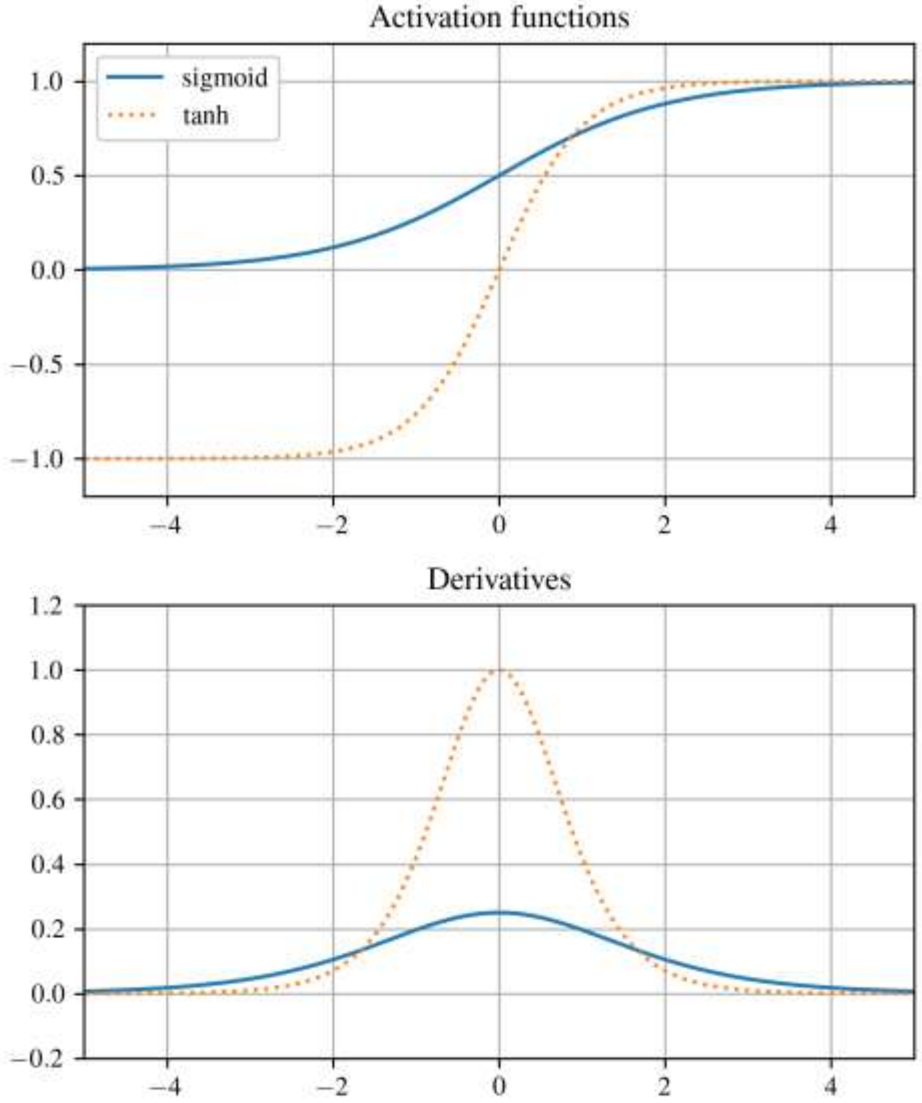

Parameter Study:

- Tanh consistently achieves lower RMSE for shallow networks
- Sigmoid requires deeper architectures to reach similar accuracy
- Both activations converge for sufficiently deep and wide models
- Activation choice primarily impacts learning efficiency, not final capability

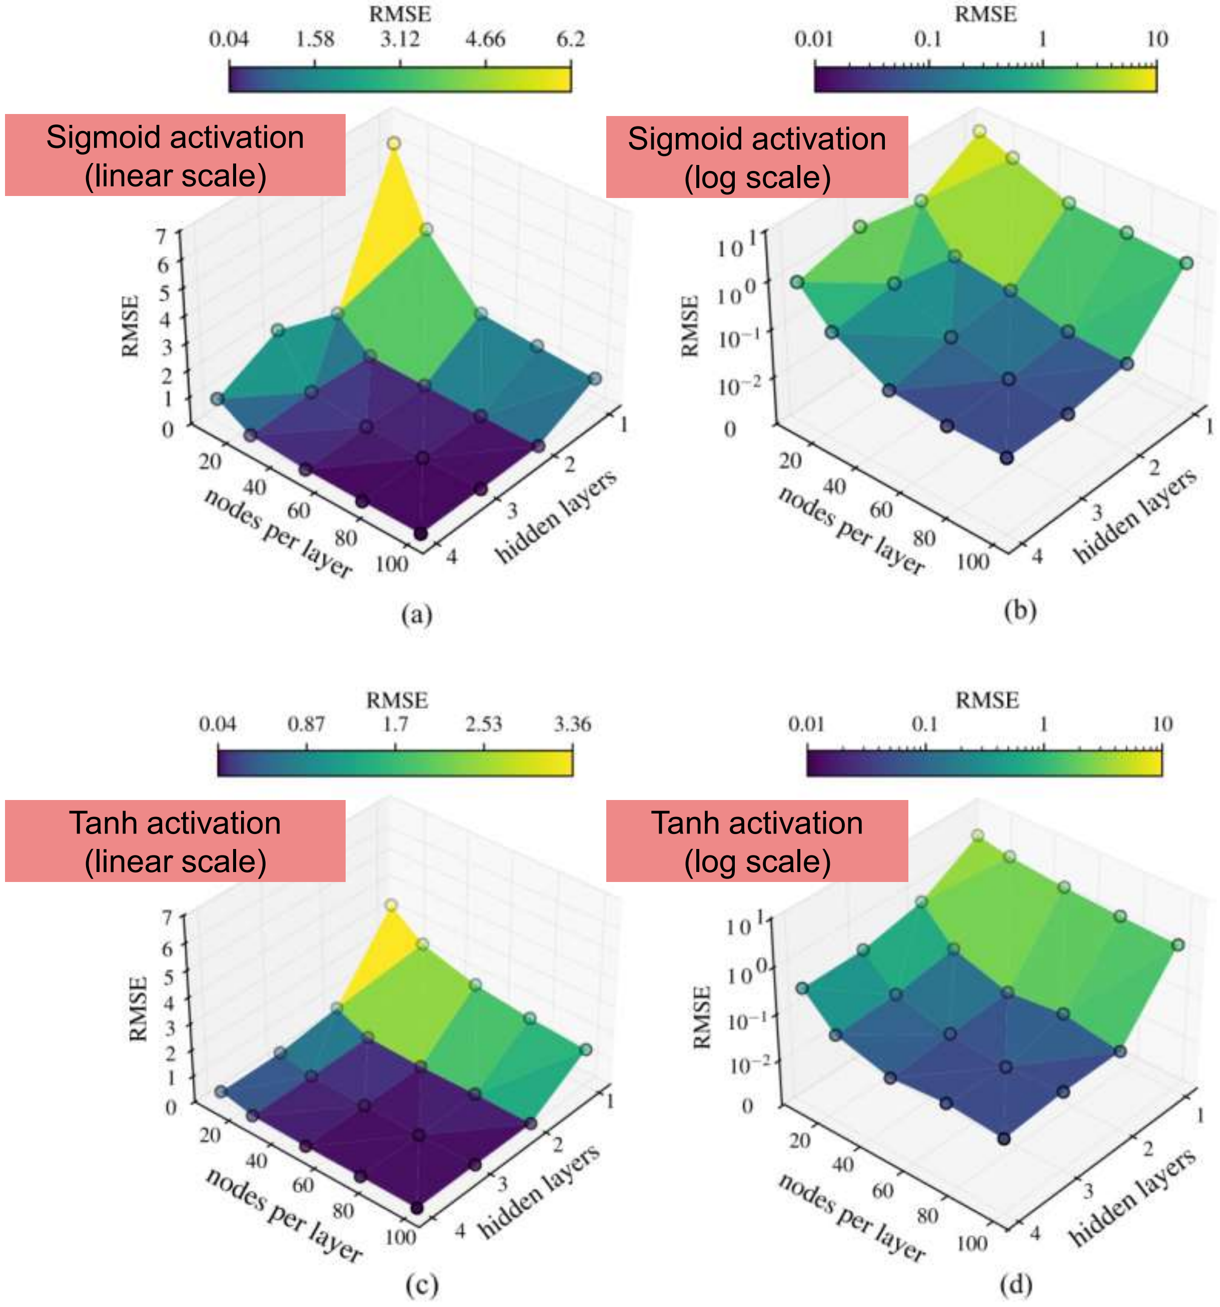

Parameter Study:

"The goal here isn't to replace existing inverse methods, but to explore what we gain—and what we give up—by enforcing physics exactly."

- Perspective:
  - Many effective methods for inverse problems in structural dynamics
  - Physics-informed machine learning is not a universal solution
  - It offers a compelling framework for combining data and physics
- Strengths of hard-constraint PIML:
  - Enables parameter estimation from indirect measurements
  - Enforces physical consistency by construction
- Limitations of hard-constraint PIML:
  - Computationally expensive due to embedded ODE solver
  - Sensitive to correctness of the assumed physical model
  - Strict enforcement can introduce high-frequency artifacts

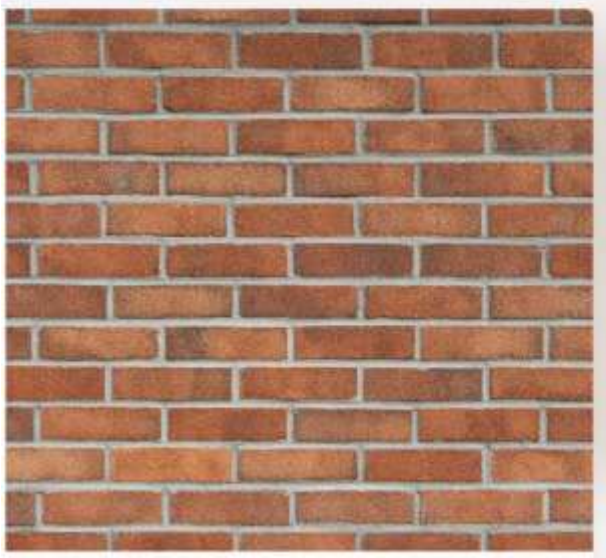

Can't go through, no matter the penalty

Acknowledgments

Air Force Office of Scientific Research

National Science Foundation

This material is based upon work supported by the National Science Foundation grant numbers CCF - 1956071, CCF-2234921, and CPS - 2237696. Additional support from the Air Force Office of Scientific Research (AFOSR) through award no. FA9550-21-1-0083. Any opinions, findings conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or the United States Air Force

Discussion

github.com/ARTS-Laboratory/Paper-2026-PIML-Part-III-Hard-Constraint-ODE-Method-for-Structural-Dynamics

Name: Austin R.J. Downey
Title: Associate Professor
Email: austindowney@gmail.com
Lab GitHub: github.com/arts-laboratory

Molinaroli College of Engineering and Computing
UNIVERSITY OF SOUTH CAROLINA