

SciML Tutorial: Learning to Model (L2M)

Thomas Beckers

Assistant Professor

Department of Computer Science

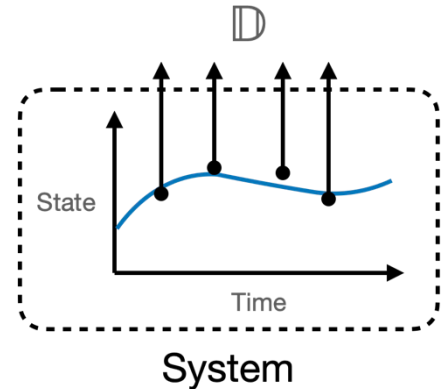
Vanderbilt University

Problem formulation

Physical system $\dot{x}(t) = \overbrace{f(x(t), u(t), t)}^{\text{unknown}}$ often Gaussian

Noisy measurements $\hat{x}_i = x(t_i) + \epsilon, \epsilon \sim \mathcal{D}$

Dataset $\mathbb{D} = \{ \{ t_k, \hat{x}_k, u_k \}_{k=0}^N \}_{i=1}^M$
Observations # Trajectories



Goal


Given the dataset \mathbb{D} , we aim to find a dynamical model that

- represents a large class of systems -> **expressive**
- **Integrates physical prior knowledge** -> **generalizes well**
- allows uncertainty quantification -> **reliable**
- accelerates simulation -> **fast solution**

How to include physics?

Observational bias


Introducing data that embody underlying physics



Symmetry

Learning bias

Inference/learning algorithm selection, loss function, optimization constraints



$$Loss = \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} + \frac{\partial w}{\partial z}$$

Inductive bias

Incorporate prior assumptions and physical constraints



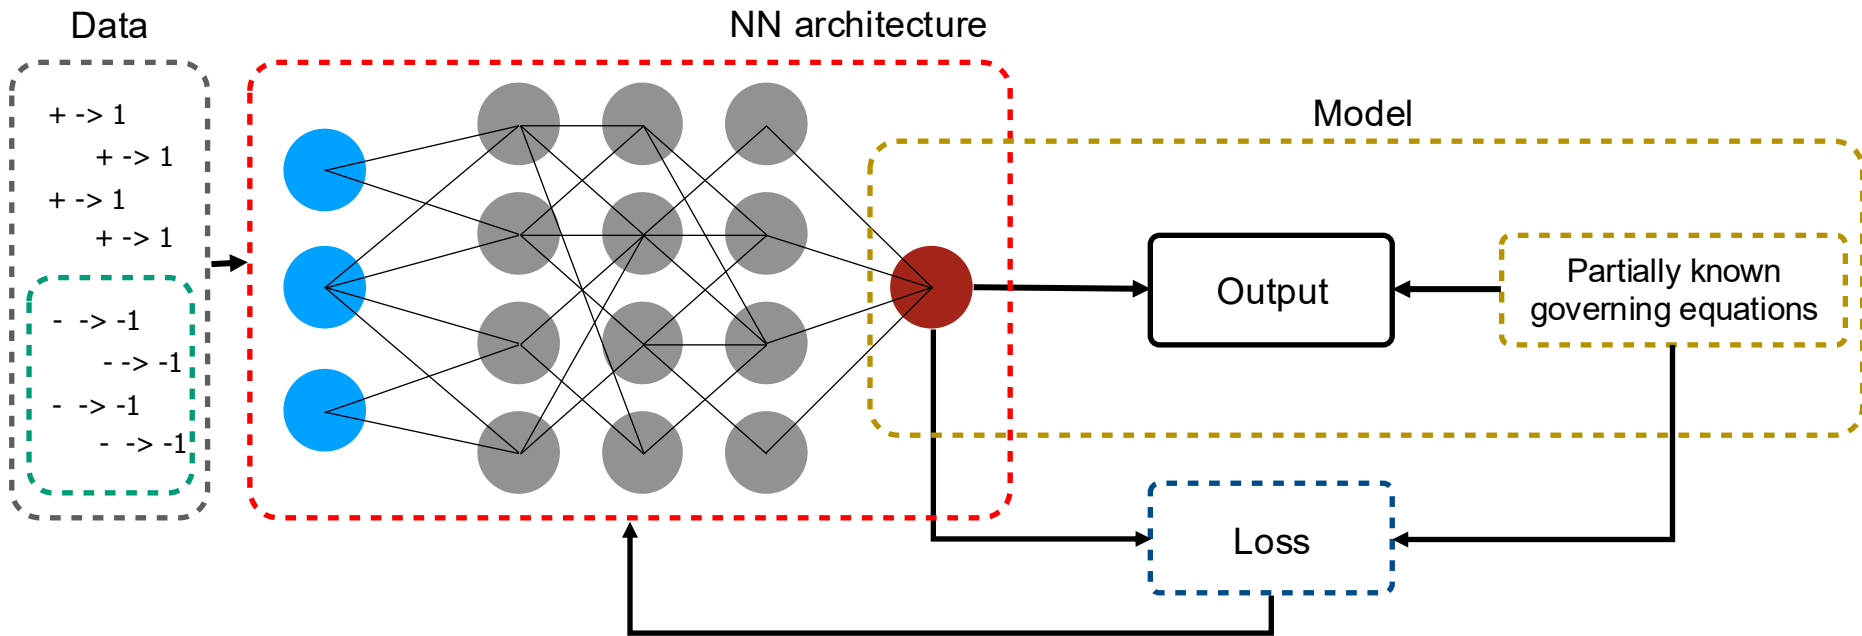
Conservation laws





Discrepancy bias

Including “terms” from partial knowledge of the model

$$\dot{x} = Ax + Bu + f(x, u)$$


Example



-  Observation bias: Introducing data that embodies underlying physics
-  Inductive bias: Prior assumptions and physical constraints
-  Learning bias: Physics based loss function, optimization algorithm
-  Discrepancy bias: Introducing data that embodies underlying physics

Combining the best of both worlds

Learning-based




Scientific principles / physics

$$\dot{x} = f(t, x, u) = \dots$$

$$0 = f\left(\frac{\partial x}{\partial t}, \frac{\partial x}{\partial z}, \dots\right) = \dots$$

physics



Structure to **constrain/inform** the model output to be **physically consistent**

Data...101011011

"Best" Learning technique?

- Neural Networks
- GNNs, RNNs, LSTMs
- Gaussian Processes
- Operator Learning


"Best" Structure?

- ODE / PDE
- Parametric
- Hamiltonian Mechanics
- Lagrangian Mechanics
- Port-Hamiltonian Systems

Choice depends on the use-case!

Combining the best of both worlds

Learning-based




Scientific principles / physics

$$\dot{x} = f(t, x, u) = \dots$$

$$0 = f\left(\frac{\partial x}{\partial t}, \frac{\partial x}{\partial z}, \dots\right) = \dots$$

physics



Structure to **constrain/inform** the model output to be **physically consistent**

Data...101011011

"Best" Learning technique?

- **Neural Networks**
- GNNs, RNNs, LSTMs
- Gaussian Processes
- Operator Learning

"Best" Structure?

- **ODE** / PDE
- Parametric
- Hamiltonian Mechanics
- Lagrangian Mechanics
- Port-Hamiltonian Systems

Descent amount of data + no physical prior knowledge

Neural ODE

Continuous-Time Ordinary Differential Equation Modeling

Given observed state-control trajectories

$$\mathbb{D} = \left\{ \left\{ t_k, \hat{x}_k, u_k \right\}_{k=0}^N, x_0 \right\}_{i=1}^M,$$

sampled at temporal points t_k , and the initial conditions x_0 . The objective is to learn parameters θ of a continuous-time dynamical model:

$$\frac{dx}{dt} = f_{\theta}(x(t), u(t), t), \quad x(t_0) = x_0,$$

s.t. the simulated states $x(t_k)$ match the observed states \hat{x}_k on the time domain $t \in [t_0, t_K]$.

Idea

- A **neural network** learns the flow field
- Output is obtained by integrating the learned ODE

$$\frac{dx}{dt} = f_{\theta}(x(t), u(t), t), \quad x(t_0) = x_0,$$

Neural network parametrized by θ

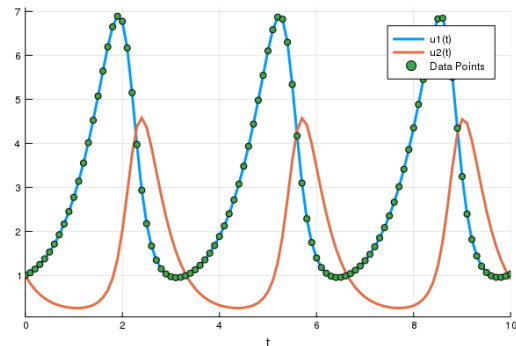

Neural ODE Loss

NODE Loss Function

Given an NODE layer $\text{ODESolve}(f_\theta(\cdot), t, x_0)$ that computes $x^i(t)$ from initial conditions x_0^i under inputs $u^i(t)$, the NODE loss measures trajectory mismatch and (optionally) model regularization:

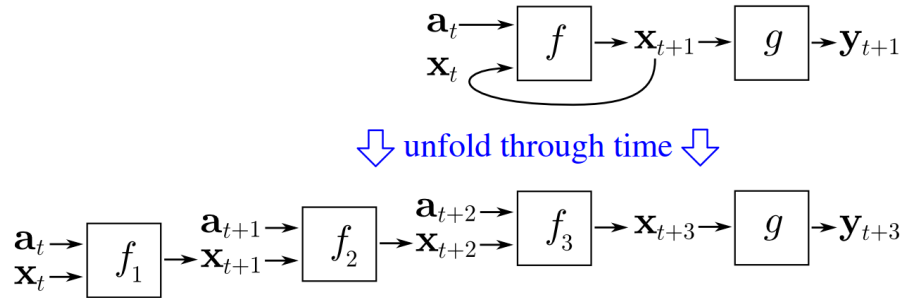
$$\mathcal{L}(\theta) = \underbrace{\frac{1}{MN} \sum_{i=1}^M \sum_{k=0}^N \|x^i(t_k) - \hat{x}_k^i\|_2^2}_{\text{trajectory data fit}} + \underbrace{\lambda \mathcal{R}(\theta)}_{\text{model regularization}} .$$

https://colab.research.google.com/github/SciML-L4DC2026/SciML-L4DC2026/blob/main/notebooks/L2M/Example_1_NODE/Example_1_NODE.ipynb




Neural ODE Training

Backpropagation through time (BPTT)



- Pros:** exact for discretized systems, numerically more stable, wide support in modern AD libraries
- Cons:** high memory cost, requires storing all intermediate variables


Adjoint sensitivity method



- Pros:** memory efficiency, no need to store intermediate variables, strong theory
- Cons:** error accumulation leading to numerical instability

Combining the best of both worlds

Learning-based




Scientific principles / physics

$$\dot{x} = f(t, x, u) = \dots$$

$$0 = f\left(\frac{\partial x}{\partial t}, \frac{\partial x}{\partial z}, \dots\right) = \dots$$

physics



Structure to **constrain/inform** the model output to be **physically consistent**

Data...101011011

"Best" Learning technique?

- Neural Networks
- GNNs, RNNs, LSTMs
- Gaussian Processes
- Operator Learning


"Best" Structure?

- ODE / PDE
- Parametric
- Hamiltonian Mechanics
- Lagrangian Mechanics
- Port-Hamiltonian Systems

Choice depends on the use-case!

Combining the best of both worlds

Learning-based




Scientific principles / physics

$$\dot{x} = f(t, x, u) = \dots$$

$$0 = f\left(\frac{\partial x}{\partial t}, \frac{\partial x}{\partial z}, \dots\right) = \dots$$

physics



Structure to **constrain/inform** the model output to be **physically consistent**

Data...101011011

"Best" Learning technique?

- Neural Networks
- GNNs, RNNs, LSTMs
- **Gaussian Processes**
- Operator Learning

"Best" Structure?

- ODE / PDE
- Parametric
- Hamiltonian Mechanics
- Lagrangian Mechanics
- **Port-Hamiltonian Systems**

Little data + Energy-based prior

Gaussian Process


Prior: Gaussian distribution over function space

$$f(x) \sim GP(m(x), k(x, x'))$$

Mean function

Covariance

$$\text{posterior} = \frac{\text{likelihood} \times \text{prior}}{\text{marginal likelihood}}$$



Gaussian Process

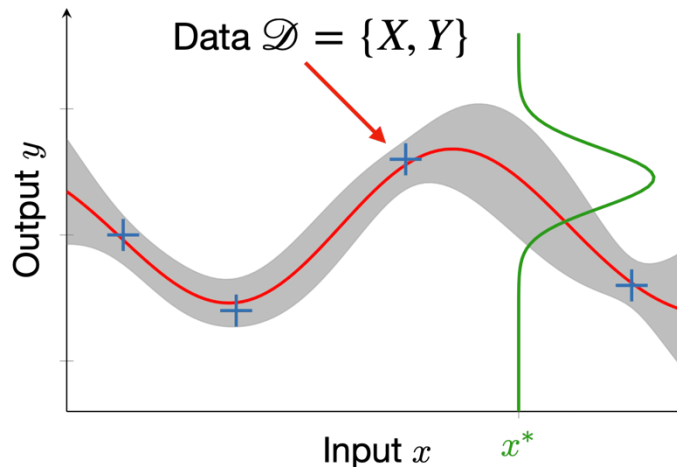
Prior: Gaussian distribution over function space

$$f(x) \sim GP(m(x), k(x, x'))$$

Mean function

Covariance

$$\text{posterior} = \frac{\text{likelihood} \times \text{prior}}{\text{marginal likelihood}}$$



Posterior

$$\mu(y^* | x^*, \mathcal{D}) = m(x^*) + k(x^*, X)[K(X, X) + \sigma_n^2 I]^{-1}(Y - m(X))$$

$$\Sigma(y^* | x^*, \mathcal{D}) = k(x^*, x^*) - k(x^*, X)[K(X, X) + \sigma_n^2 I]^{-1}k(x^*, X)^\top$$

Test input

Gaussian Process

Posterior mean

$$\mu(y | \mathbf{x}, \mathcal{D}) = m(\mathbf{x}) + \sum_i^N w_i k(\mathbf{x}, \mathbf{X}_i)$$

Data points Covariance


Prior mean function Weighting factor

Via mean function

$$m(\mathbf{x}) = \dots$$

Discrepancy bias

Via covariance function



Inductive bias


Linear operators

GPs are closed under linear operators


$$f(x) = \mathcal{L}_x g(x) \sim GP(\mathcal{L}_x \mu, \mathcal{L}_x k \mathcal{L}_x^\top)$$

Example: Differential operator $\mathcal{L}_x = \frac{\partial}{\partial x}$

$$\begin{bmatrix} y \\ Y' \end{bmatrix} = \mathcal{N} \left(0, \begin{bmatrix} k(\mathbf{x}, \mathbf{x}) & \frac{\partial}{\partial x} k(\mathbf{x}, \mathbf{X}) \\ \frac{\partial}{\partial x} k(\mathbf{x}, \mathbf{X})^\top & \frac{\partial}{\partial x} k(\mathbf{X}, \mathbf{X}) \frac{\partial}{\partial x} \end{bmatrix} \right)$$



Physics structure



[A. Van Der Schaft and D. Jeltsema]

Port-Hamiltonian system

Interconnection Dissipation Hamiltonian

$$\dot{x} = [J(x) - R(x)] \frac{\partial H}{\partial x} + G(x)u$$

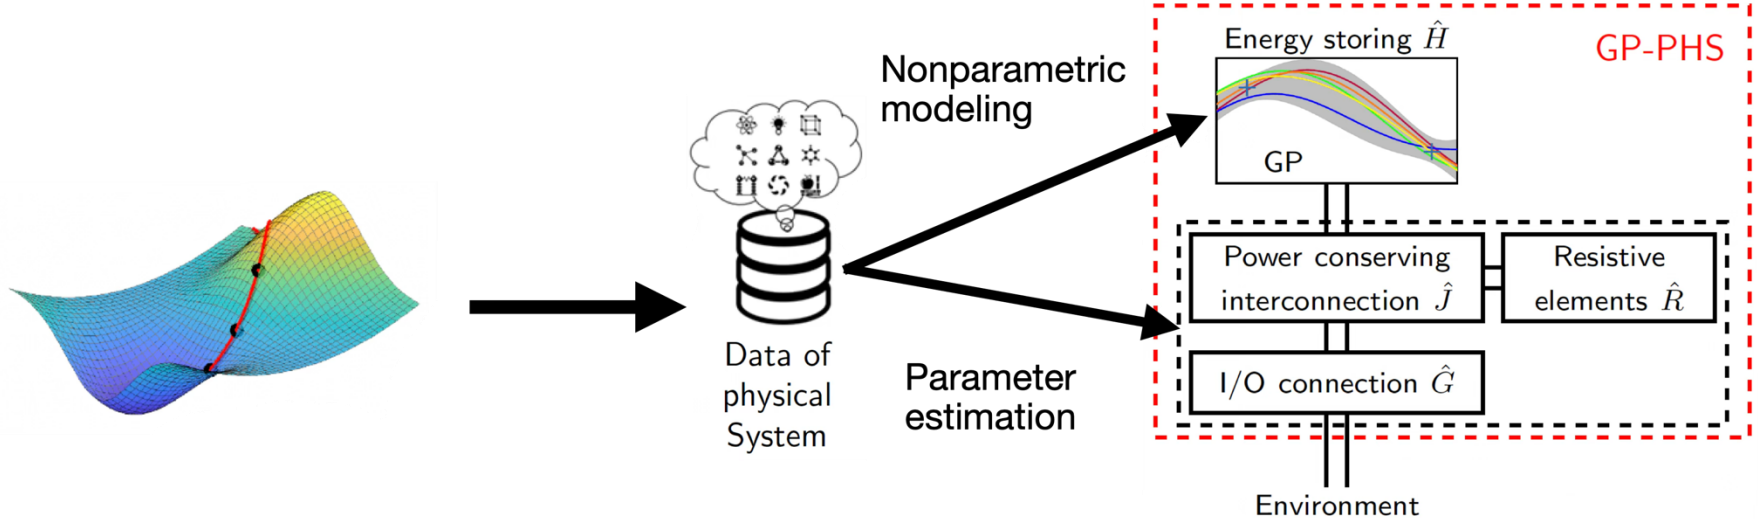
$$y = G(x)^\top \frac{\partial H}{\partial x}$$

Input

Output

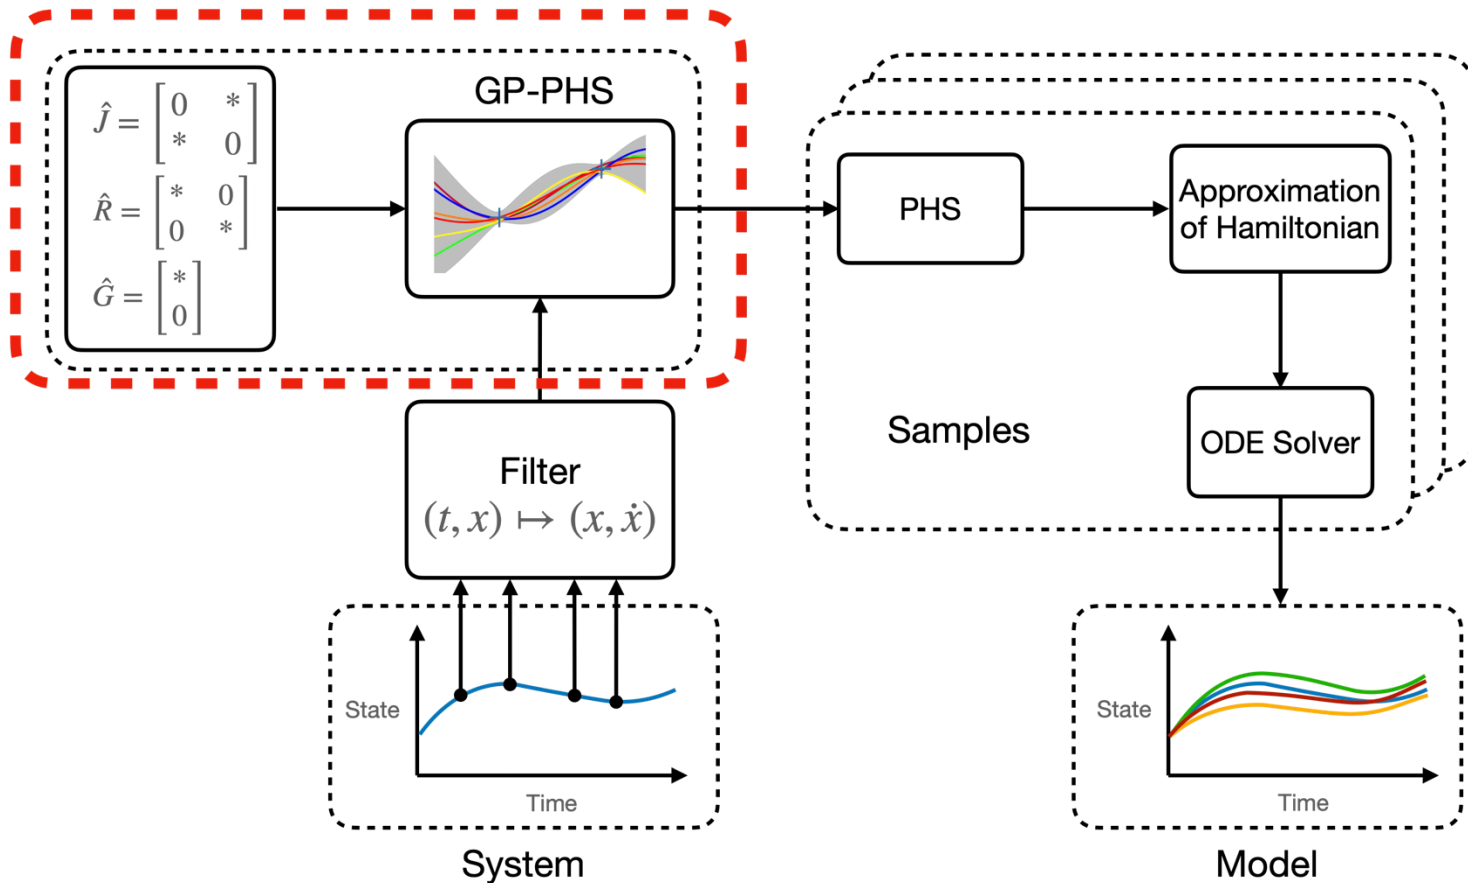
- Skew-symmetric J ensures **energy conservation**
- Connection with another PHS **preserves the structure**
- Input and output defines a **passive systems**
- Suitable for **multi-domain applications** via energy flows

GP-PHS



- Physics-constrained Gaussian Process model
- GP-PHS includes **all possible PHS** under the GP prior on the Hamiltonian

GP-PHS



GP-PHS: Training

GP prior on Hamiltonian

$$\dot{x} = [J(x) - R(x)] \frac{\partial H}{\partial x} + G(x)u \quad \longrightarrow \quad \dot{x} \sim GP(\hat{G}(x)u, k_{phs}(x, x'))$$

Linear operator

Port-Hamiltonian Kernel

$$k_{phs}(x, x') = \sigma_f^2 [\hat{J}(x) - \hat{R}(x)] \left[\frac{\partial}{\partial x \partial x'} \underbrace{\exp(-\|x - x'\|_\Lambda^2)}_{\text{Squared exp. kernel}} \right] \overbrace{[\hat{J}(x') - \hat{R}(x')]^\top}^{\text{Parametrized estimates}}$$

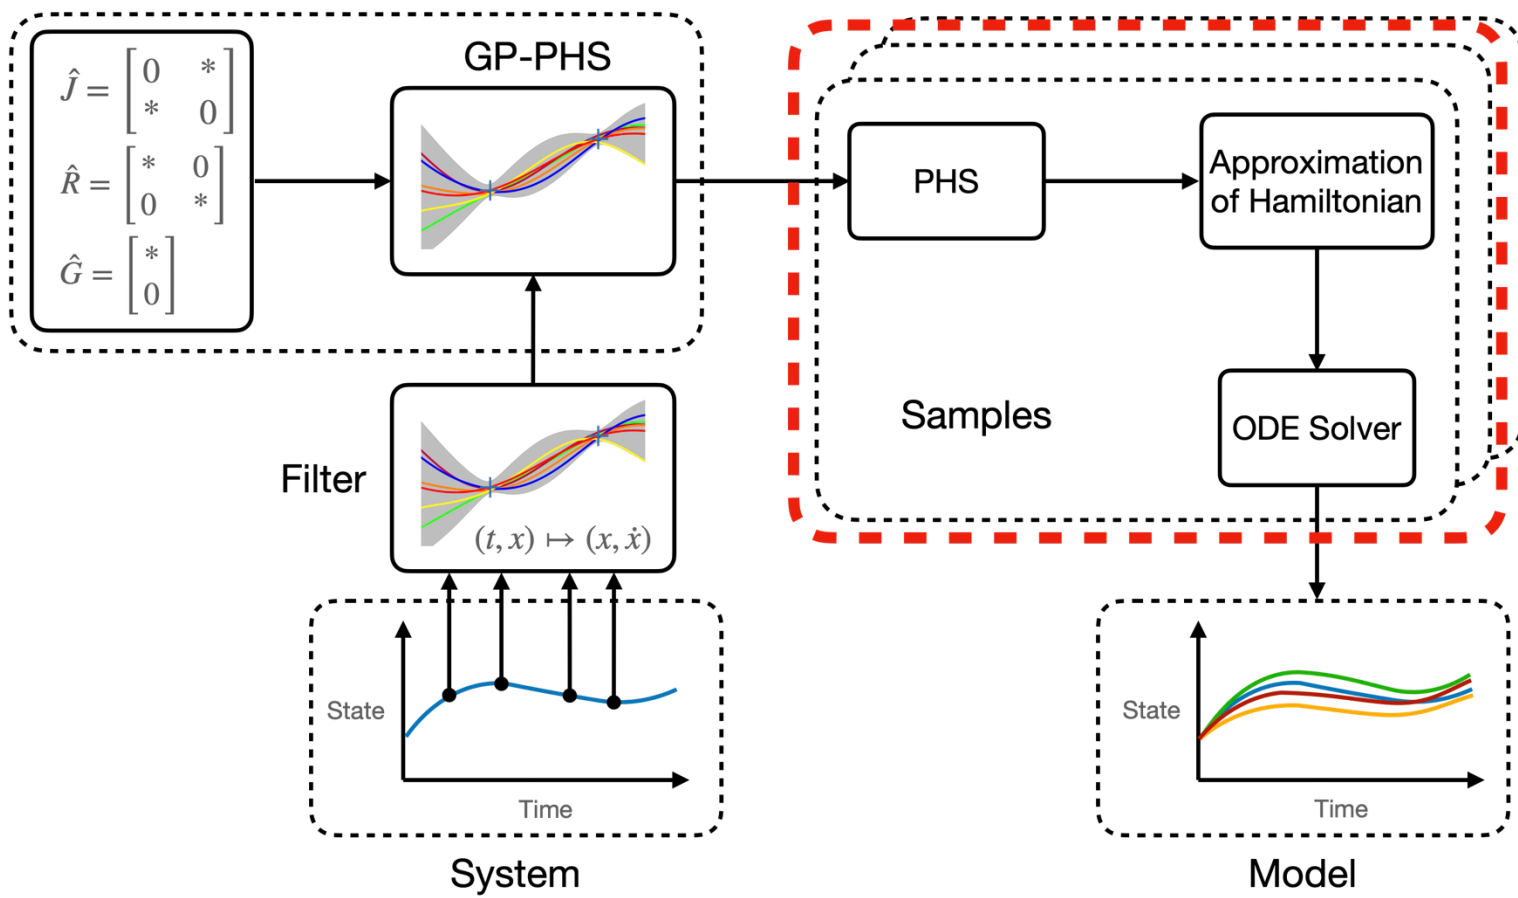
Optimization problem

$$\min_{\theta} -\log p(\dot{X} \mid \theta, X) \sim \dot{X}^\top [K_{phs} + \Delta]^{-1} \dot{X} + \log |K_{phs} + \Delta|$$

θ = Kernel parameters + unknown PHS parameters

$$\text{s.t. } J = -J^\top, \quad R \succeq 0$$

GP-PHS



GP-PHS: Sampling

PHS

$$\begin{bmatrix} \dot{X} \\ \hat{f}(x^*) \end{bmatrix} = \mathcal{N} \left(0, \begin{bmatrix} K_{phs} & k_{phs}(X, x^*) \\ k_{phs}(X, x^*)^\top & k_{phs}(x^*, x^*) \end{bmatrix} \right)$$

$$\dot{x} = \hat{f}(x, \omega)$$

↑
Not callable
(computational expensive)

Approximation
of Hamiltonian

$$\begin{bmatrix} \dot{X} \\ \hat{H}(x^*) \end{bmatrix} = \mathcal{N} \left(0, \begin{bmatrix} K_{phs} & k_{\dot{x}H}(X, x^*) \\ k_{\dot{x}H}(X, x^*)^\top & k_{HH}(x^*, x^*) \end{bmatrix} \right)$$

$$\hat{H}(x^*, \omega) \xrightarrow{\text{Approx.}} \hat{H}^*(x^*)$$

ODE Solver

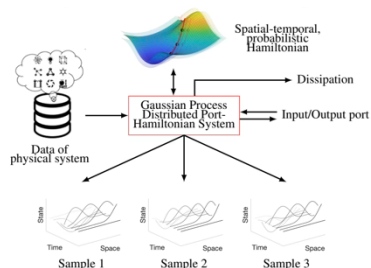
$$\dot{x} = [\hat{J}(x) - \hat{R}(x)] \frac{\partial \hat{H}^*}{\partial x} + \hat{G}(x)u \xrightarrow{\text{Solver}} x(t)$$

Approximation of H: Callable and structure preserving

Use-cases and extensions


PDEs

Energy-based learning of PDE dynamics




Switching systems

Learning the switching surface



Control

Passivity-based control approaches



...and more

Full Bayesian: Learning from input-output data

Ewering, Jan-Hendrik, et al. "Learning Dynamics from Input-Output Data with Hamiltonian Gaussian Processes." L4DC, 2026

Finite-element framework

Courteville, Florian, et al. "PFEM-GP-dpHS: a finite element framework for combining Gaussian processes and infinite-dimensional port-Hamiltonian systems." L4DC, 2026

GP-PHS: Implementation

```
CO Gaussian Process Port-Hamiltonian Systems ☆ ☁
File Edit View Insert Runtime Tools Help
Q Commands | + Code + Text | ▶ Run all ▼

[ ] !pip install gpytorch
    !pip install torchdiffeq

Show hidden output

[ ] import torch
    import gpytorch
    import math
    import numpy as np
    from torchdiffeq import odeint
    from matplotlib import pyplot as plt
    from scipy.interpolate import RegularGridInterpolator
```

https://colab.research.google.com/github/SciML-L4DC2026/SciML-L4DC2026/blob/main/notebooks/L2M/Example_2_GPPHS/GP_PHS.ipynb

Implementation based on Torch and GPyTorch

- Use of Torch features such as GPU learning
- Tools from GPyTorch (speed up computation, etc.)
- PHS-Kernel implemented based on

RBFKernelGradGrad

```
class gpytorch.kernels.RBFKernelGradGrad(ard_num_dims=None, batch_shape=None, active_dims=None,
lengthscale_prior=None, lengthscale_constraint=None, eps=1e-06, **kwargs) [source]
```

Computes a covariance matrix of the RBF kernel that models the covariance between the values and first and second (non-mixed) partial derivatives for inputs \mathbf{x}_1 and \mathbf{x}_2 .

See `gpytorch.kernels.Kernel` for descriptions of the lengthscale options.

$$\begin{bmatrix} k(\mathbf{x}, \mathbf{x}) & \frac{\partial}{\partial \mathbf{x}} k(\mathbf{x}, \mathbf{X}) \\ \frac{\partial}{\partial \mathbf{x}} k(\mathbf{x}, \mathbf{X})^\top & \frac{\partial}{\partial \mathbf{x}} \frac{\partial}{\partial \mathbf{x}} k(\mathbf{X}, \mathbf{X}) \frac{\partial}{\partial \mathbf{x}} \end{bmatrix}$$

Learning to Control (L2M) Summary

- **L2M via data + structure/physics**

- Combining data and physics to improve model efficiency and generalization
- Embedding prior knowledge through model structure, constraints, and inductive biases
- Quantifying uncertainty to assess model confidence and guide decision-making

- **Challenges**

- Balancing model expressiveness and physical interpretability
- Efficient training and uncertainty-aware prediction
- Learning from sparse, heterogeneous, noisy, and incomplete observations

- **Code**

- github.com/pnnl/neuromancer
- www.tbeckers.com

