

# Embodied Vision

Physical Modeling: Learning-based Dynamics

Tsung-Wei Ke

Spring 2025



# Physical Modeling We Have Covered

- We have covered mass-spring system, position-based dynamics, incremental potential contact and material point methods
- Pros:
  - Physically accurate if parameters are well tuned
- Cons:
  - Need to manually select a lot of parameters. For instance, connection pattern and stiffness for mass-spring system; material type, boundary conditions,  $E$ ,  $\nu$  for MPM...
  - These frameworks lack of generalization, each of which is only good for certain scenarios

# Learn Physically Accurate Dynamics without the Need for Parameter Tuning but Generalize Well

- Ideas:
  - Pure data-driven approaches: Fit models on large-scale data
  - Hybrid approaches: Data-driven supervision and physic-based regularization


# Content

- Pure Data-driven Approaches
- Hybrid Approaches
  - Physics-Informed Diffusion Model
  - Physics-Informed Neural Network
  - Neural ODE
  - Lagrangian Network / Hamiltonian Network
  - Neural Operator
  - Fourier Neural Operator

# Content


- Pure Data-driven Approaches
- Hybrid Approaches
  - Physics-Informed Diffusion Model
  - Physics-Informed Neural Network
  - Neural ODE
  - Lagrangian Network / Hamiltonian Network
  - Neural Operator
  - Fourier Neural Operator

# Pure Data-Driven Approaches



- All you need is to collect a large-scale dataset, and fit generative models on the dataset

# Generative Modeling for Simulation



- What are the input representations?
- What are the model architectures?
- What are the output representations?
- How to obtain ground truths?

---

# Learning to Simulate Complex Physics with Graph Networks


---

**Alvaro Sanchez-Gonzalez**<sup>\*1</sup> **Jonathan Godwin**<sup>\*1</sup> **Tobias Pfaff**<sup>\*1</sup> **Rex Ying**<sup>\*12</sup> **Jure Leskovec**<sup>2</sup>  
**Peter W. Battaglia**<sup>1</sup>


- What are the input representations?
  - Particles
- What are the model architectures?
  - Graph neural networks
- What are the output representations?
  - Particle position / velocity ...
- How to obtain ground truths?
  - Generate with PBD, MPM, ...

# Step 1: Encode a Graph of Particles


(a)  $X^{t_0}$



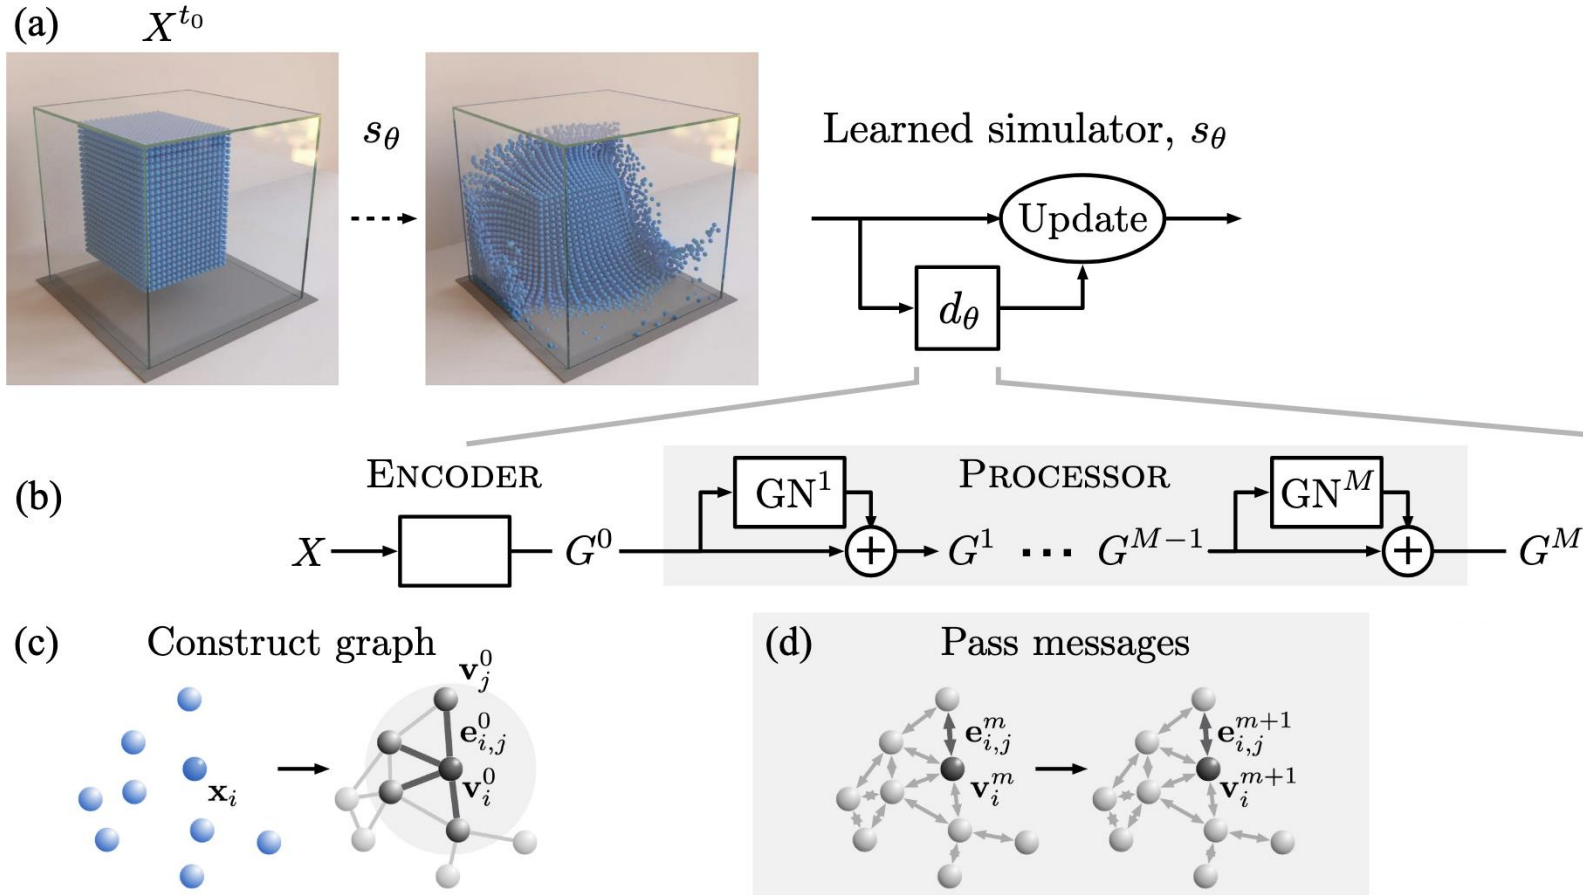
(b)



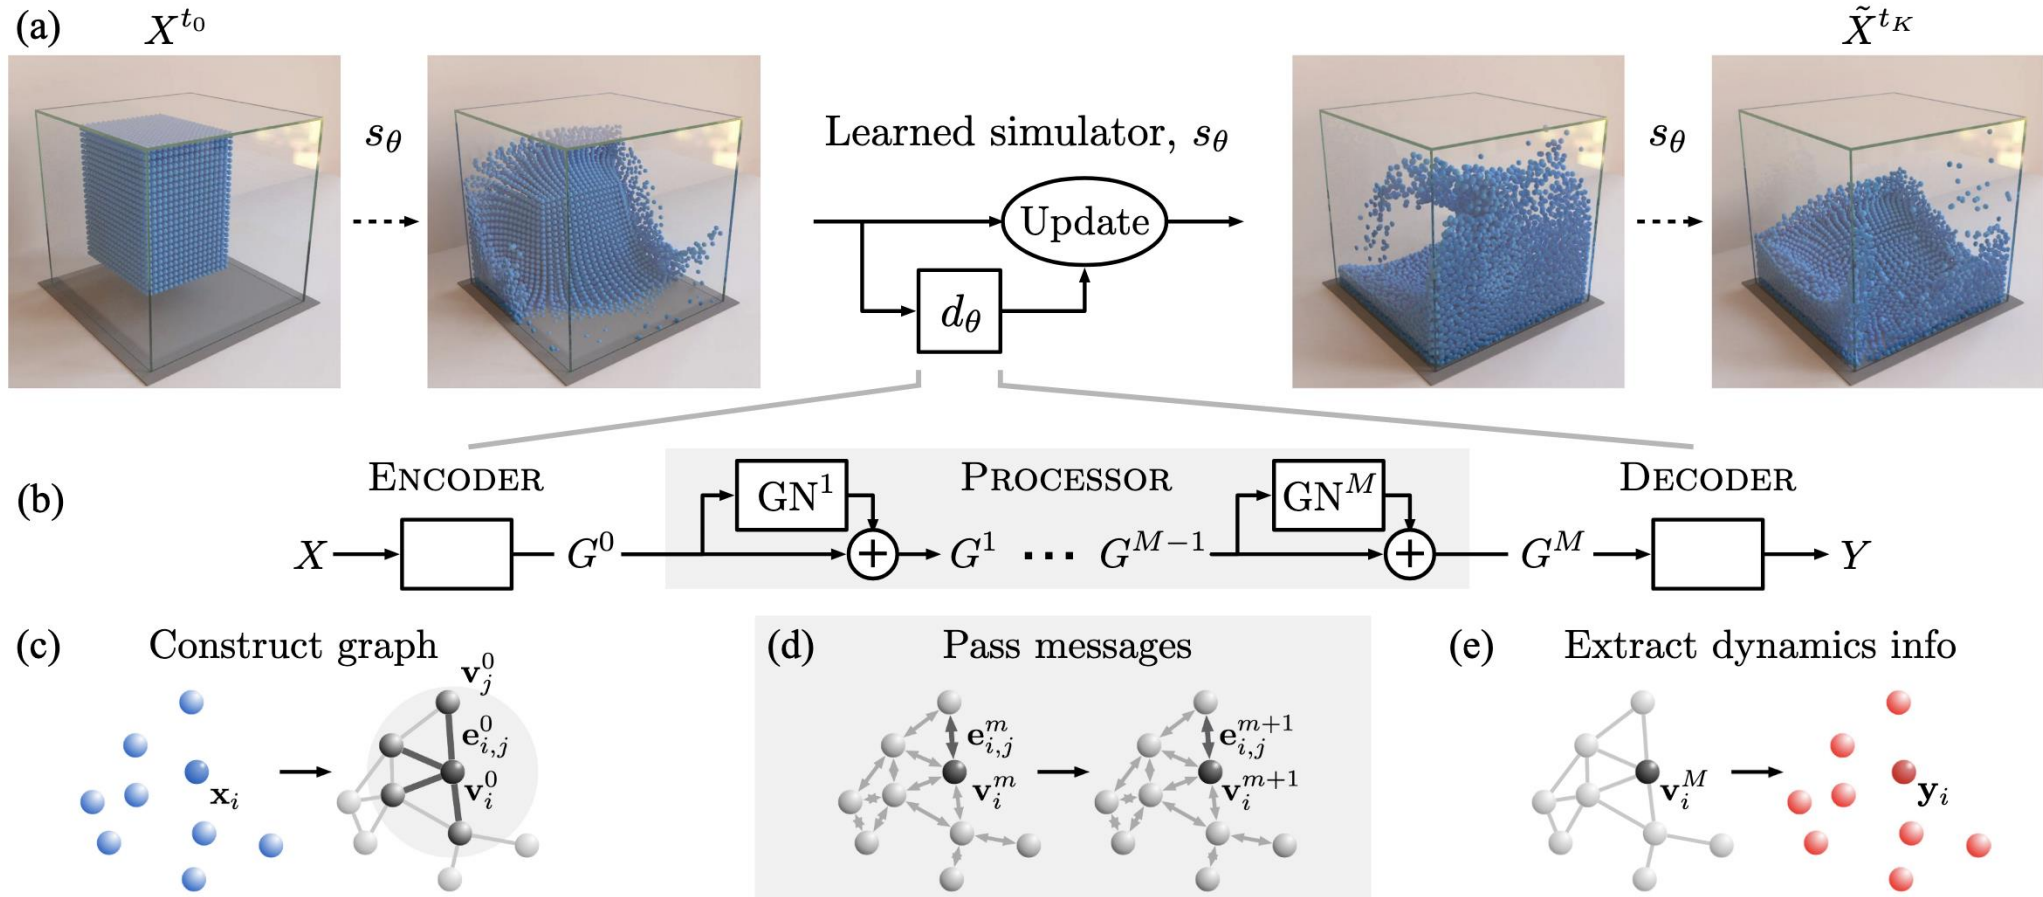
(c)




# Step 2: Contextualize the Graph with Graph Neural Network



# Step 3: Decode Particle Motion



# Results




# Learning 3D Particle-based Simulators from RGB-D Videos



**William F. Whitney\***, **Tatiana Lopez-Guevara\***, **Tobias Pfaff**, **Yulia Rubanova**,  
**Thomas Kipf**, **Kimberly Stachenfeld**, **Kelsey R. Allen**  
Google DeepMind

- What are the input representations?
  - Pixel particles
- What are the model architectures?
  - Graph neural networks
- What are the output representations?
  - Particle position / velocity ...
- How to obtain ground truths?
  - RGB-D videos



# Similar Frameworks Supervised with Rendering Loss



# Results




# Results




# Results: Scene Editing

Origin scene



Delete cylinder




Delete floor



# However, Pure Data-Driven Approaches May Not Work We Learn the Same Lesson from Video Generative Models



Video generated by Veo3




Video generated by Veo3

# Content

- Pure Data-driven Approaches
- Hybrid Approaches
  - Physics-Informed Diffusion Model
  - Neural ODE
  - Lagrangian Network / Hamiltonian Network
  - Physics-Informed Neural Network
  - Neural Operator
  - Fourier Neural Operator

# Physic-based Regularization



# Content

- Pure Data-driven Approaches
- Hybrid Approaches
  - Physics-Informed Diffusion Model
  - Physics-Informed Neural Network
  - Neural ODE
  - Lagrangian Network / Hamiltonian Network
  - Neural Operator
  - Fourier Neural Operator

# PHYSICS-INFORMED DIFFUSION MODELS

## **Jan-Hendrik Bastek**

Dept. of Mechanical and Process Eng.  
ETH Zurich  
Zurich, Switzerland  
jbastek@ethz.ch

## **WaiChing Sun**

Dept. of Civil Eng. and Eng. Mechanics  
Columbia University  
New York, NY, USA  
wsun@columbia.edu

## **Dennis M. Kochmann**

Dept. of Mechanical and Process Eng.  
ETH Zurich  
Zurich, Switzerland  
dmk@ethz.ch

- Enforce samples generated from the learned distribution conform to the known physical constraints

# We Can Formulate Physic Laws as a set of PDEs


- Constraints:

$$\mathcal{F}[\mathbf{u}(\boldsymbol{\xi})] = \mathbf{0}, \quad \boldsymbol{\xi} = (\xi_1, \xi_2, \dots, \xi_d)^\top \in \Omega, \quad \mathbf{u} = (u_1(\boldsymbol{\xi}), u_2(\boldsymbol{\xi}), \dots, u_c(\boldsymbol{\xi}))^\top \in \mathbb{R}^c,$$

where  $\mathcal{F}$  denote differential operators,  $\boldsymbol{\xi}$  denote elements in domain  $\Omega$  and  $\mathbf{u}$  denote the solution field satisfies the set of PDES

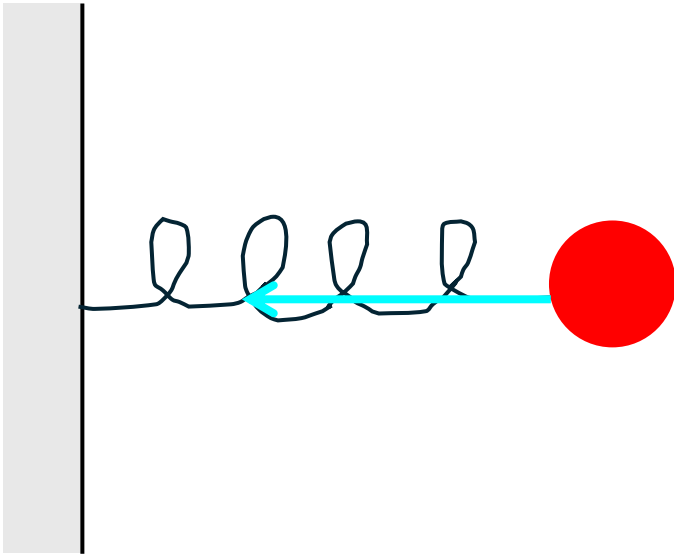
- Boundary conditions:  $\mathcal{B}[\mathbf{u}(\boldsymbol{\xi})] = \mathbf{0}, \quad \boldsymbol{\xi} \in \partial\Omega$
- We can compose both factors into a residual operator:  $\mathcal{R}(\mathbf{x}_0) := \begin{bmatrix} \mathcal{F}[\mathbf{x}_0] \\ \mathcal{B}[\mathbf{x}_0] \end{bmatrix}$
- To enforce physical constraints, we minimize the residual  $\mathcal{R}(\mathbf{x}_0)$  of sample  $\mathbf{x}_0$

# Example: Position Points on a Circle



- Boundary conditions:  
$$(x - x_c)^2 + (y - y_c)^2 = r^2$$

# Example: Spring Motion



- Physic law:

$$m\ddot{x} = -kx$$

- Boundary conditions:

$$(x - x_c)^2 + (y - y_c)^2 = r^2$$

# Integrate into the Diffusion Model

- To maintain the probabilistic perspective of generative models, we can introduce the residuals as *virtual observables*

$$q_{\mathcal{R}}(\hat{\mathbf{r}}|\mathbf{x}_0) = \mathcal{N}(\hat{\mathbf{r}}; \mathcal{R}(\mathbf{x}_0), \sigma^2 \mathbf{I})$$

clean data sample

# Integrate into the Diffusion Model

- To maintain the probabilistic perspective of generative models, we can introduce the residuals as *virtual observables*

$$q_{\mathcal{R}}(\hat{\mathbf{r}}|\mathbf{x}_0) = \mathcal{N}(\hat{\mathbf{r}}; \mathcal{R}(\mathbf{x}_0), \sigma^2 \mathbf{I})$$

- The likelihood of *virtual observables*:


$$p_{\theta}(\hat{\mathbf{r}}) = \int p_{\theta}(\hat{\mathbf{r}}, \mathbf{x}_0) d\mathbf{x}_0 = \int q_{\mathcal{R}}(\hat{\mathbf{r}}|\mathbf{x}_0) p_{\theta}(\mathbf{x}_0) d\mathbf{x}_0 = \mathbb{E}_{\mathbf{x}_0 \sim p_{\theta}(\mathbf{x}_0)} q_{\mathcal{R}}(\hat{\mathbf{r}}|\mathbf{x}_0).$$

- Generating a physically correct sample is to maximize the likelihood of *virtual observables*:

$$\arg \max_{\theta} \mathbb{E}_{\hat{\mathbf{r}}} [\log p_{\theta}(\hat{\mathbf{r}})] = \arg \max_{\theta} \mathbb{E}_{\mathbf{x}_0 \sim p_{\theta}(\mathbf{x}_0)} [\log q_{\mathcal{R}}(\hat{\mathbf{r}} = \mathbf{0}|\mathbf{x}_0)].$$

# Step 1: Estimate $\mathbf{x}_0^*$

- Idea 1: accelerate sampling with DDIM

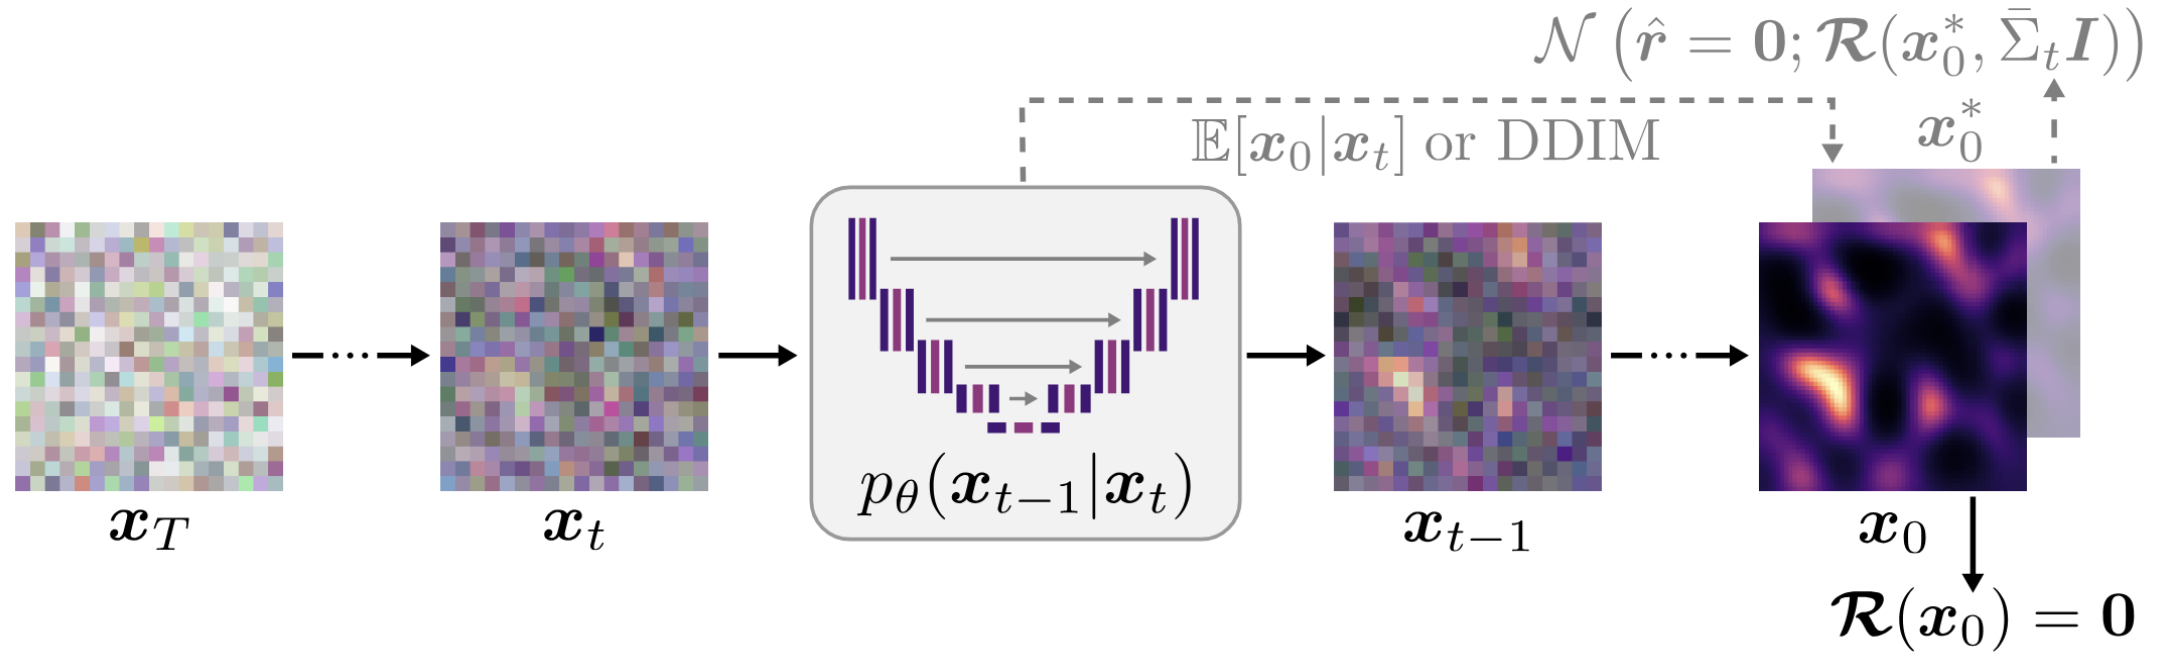


- Idea 2: estimate directly from  $\mathbf{x}_t$

$$q(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t}\mathbf{x}_0, (1 - \bar{\alpha}_t)\mathbf{I})$$

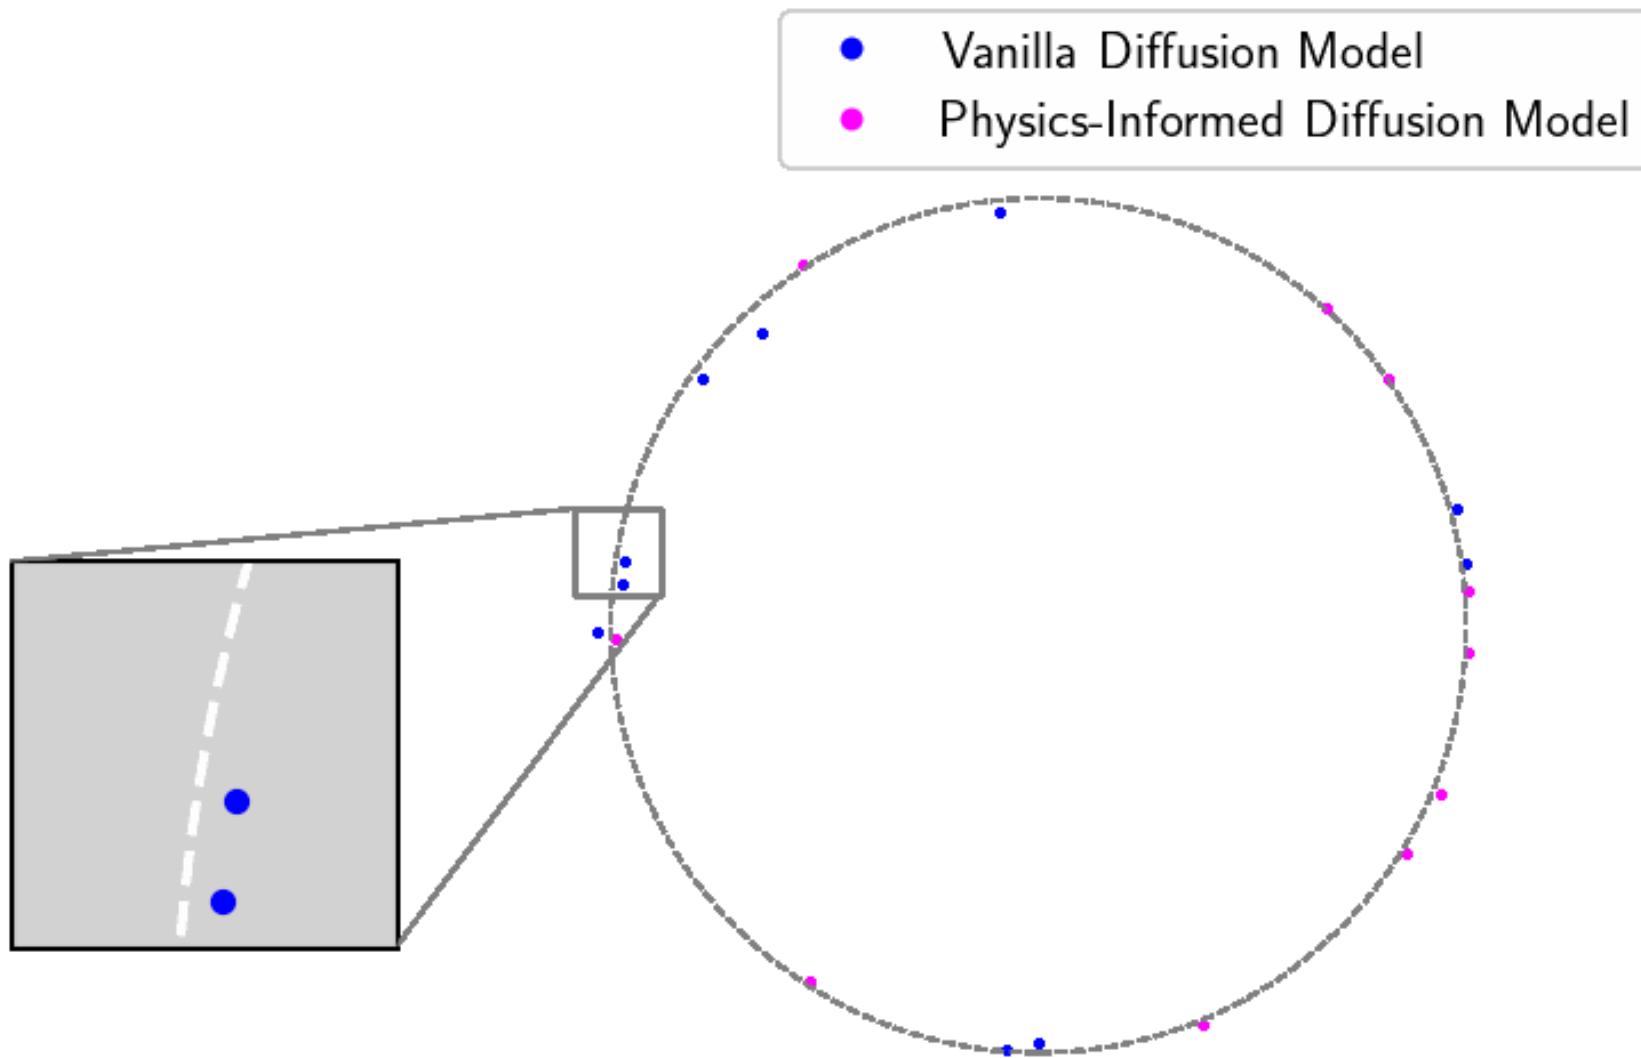
$$\mathbf{x}_0 \approx \hat{\mathbf{x}}_0 = (\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t}\epsilon_\theta(\mathbf{x}_t)) / \sqrt{\bar{\alpha}_t}$$

# Step 2: Minimize Residual Loss

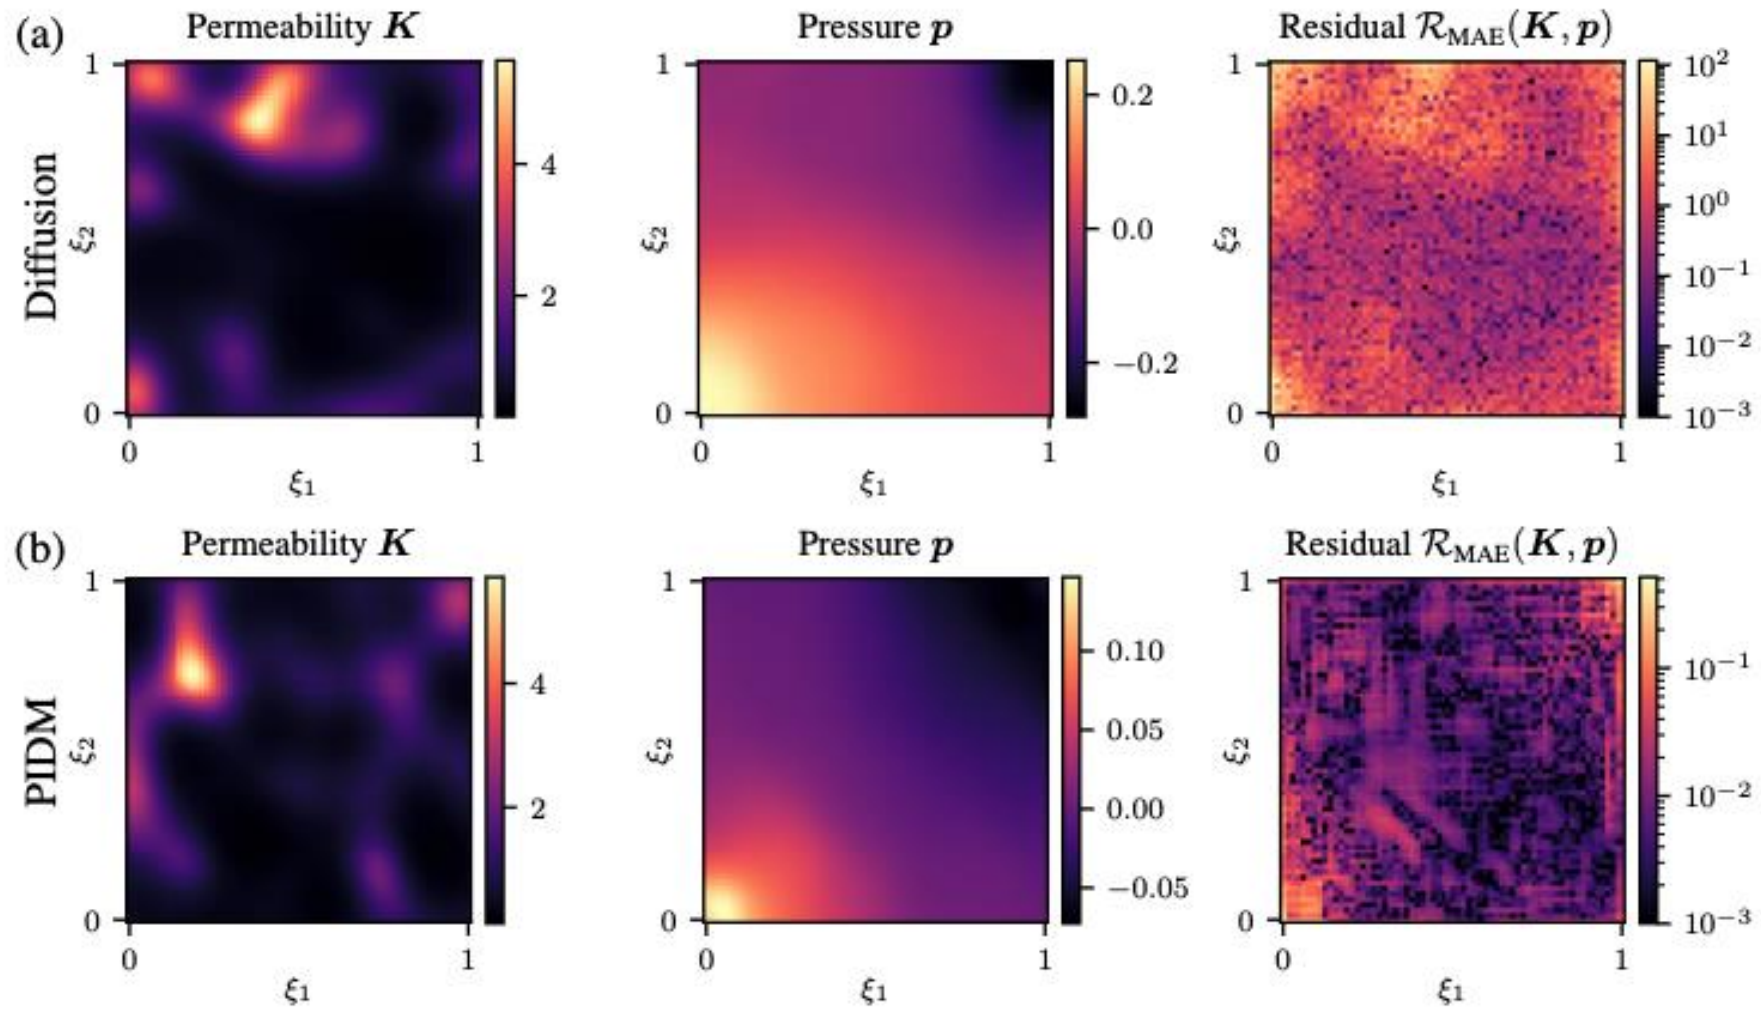


$$L_{\text{PIDM}}(\theta) = \mathbb{E}_{t \sim [1, T], \mathbf{x}_{0:T} \sim q(\mathbf{x}_{0:T})} \left[ \lambda_t \|\mathbf{x}_0 - \hat{\mathbf{x}}_0(\mathbf{x}_t, t)\|^2 + \frac{1}{2\bar{\Sigma}_t} \|\mathcal{R}(\mathbf{x}_0^*(\mathbf{x}_t, t))\|^2 \right],$$

# Results



# Results



# Content

- Pure Data-driven Approaches
- Hybrid Approaches
  - Physics-Informed Diffusion Model
  - Physics-Informed Neural Network
  - Neural ODE
  - Lagrangian Network / Hamiltonian Network
  - Neural Operator
  - Fourier Neural Operator

# A More General Formulation of Physics-Informed Neural Network (PINN)

- Core idea of physics-informed diffusion model is to impose constraints based on physic laws on the model output


- A general form of PDEs:

$$\mathbf{u}_t + \mathcal{N}[\mathbf{u}] = \mathbf{0}, \mathbf{x} \in \Omega, t \in [0, T],$$

where  $\mathbf{u}$  is the solution,  $\mathcal{N}[\cdot]$  denotes a non-linear differential operator.

- Idea: learn a neural network predicting solutions of PDEs to bypass time-consuming numerical simulation

# PINN



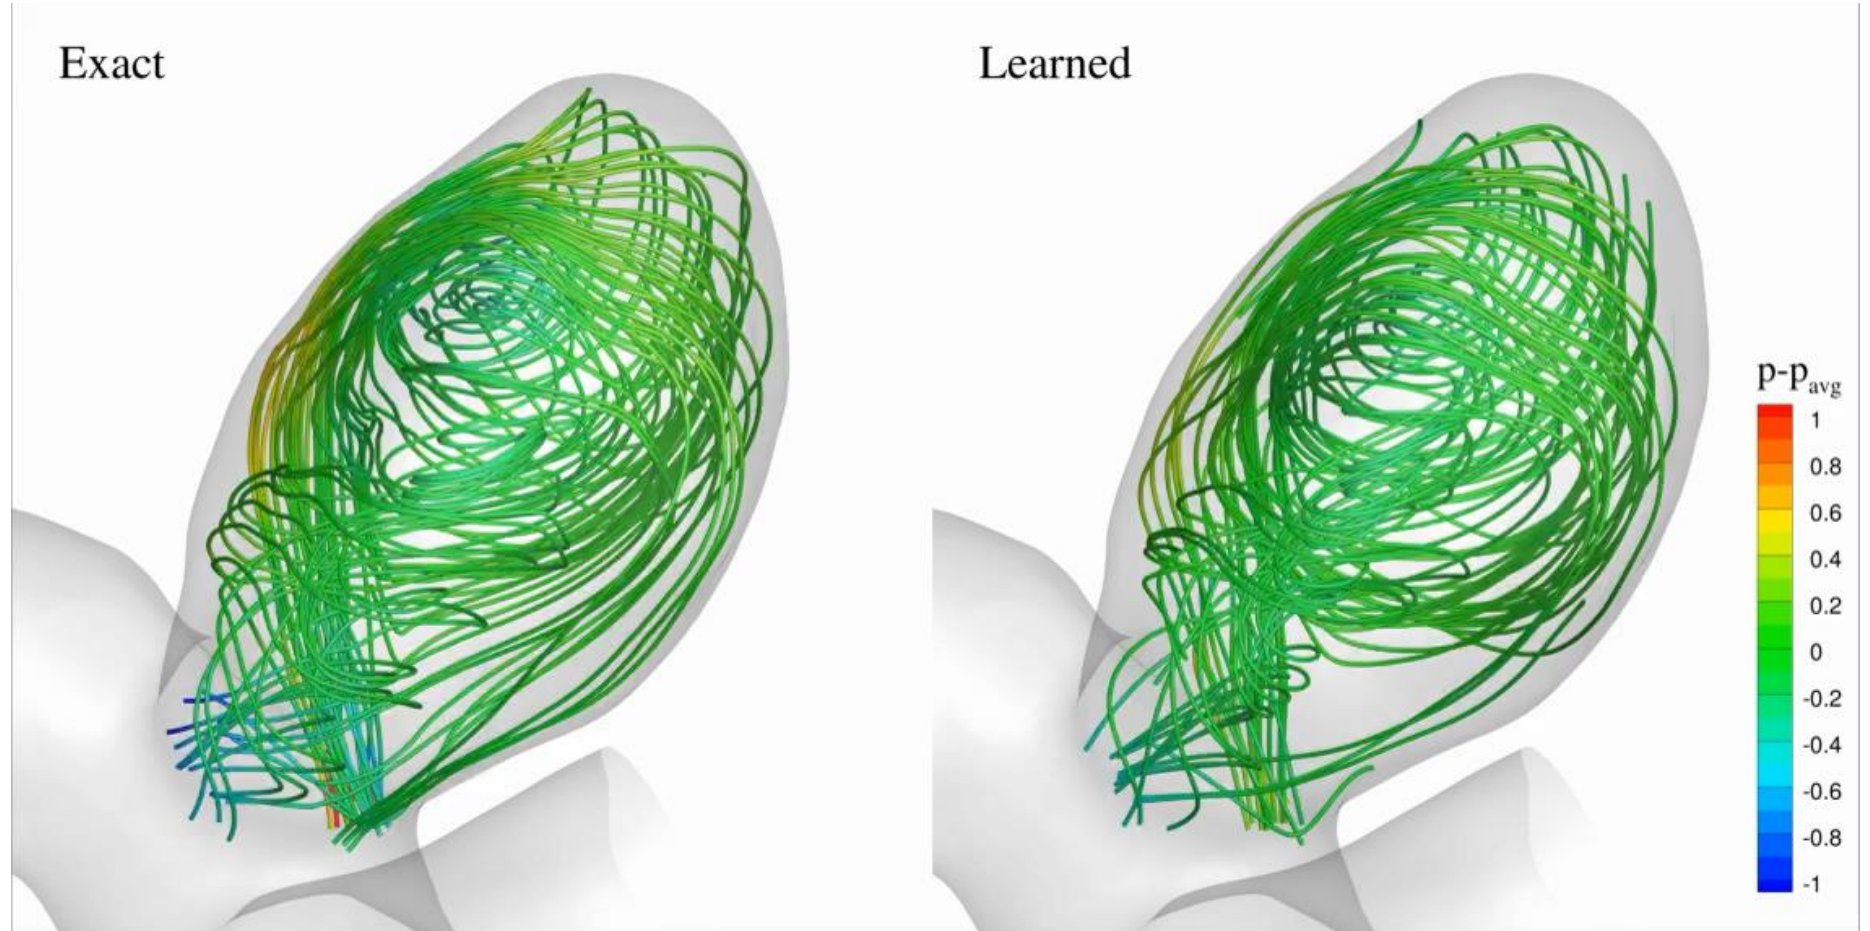
## Naviers-Stokes loss

$$\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} + \frac{\partial w}{\partial z} - 0$$
$$\frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} + v \frac{\partial u}{\partial y} + w \frac{\partial u}{\partial z} + \frac{\partial P}{\partial x} - \frac{1}{Re} \left( \frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} + \frac{\partial^2 u}{\partial z^2} \right) = 0$$
$$\frac{\partial v}{\partial t} + u \frac{\partial v}{\partial x} + v \frac{\partial v}{\partial y} + w \frac{\partial v}{\partial z} + \frac{\partial P}{\partial y} - \frac{1}{Re} \left( \frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial y^2} + \frac{\partial^2 v}{\partial z^2} \right) = 0$$
$$\frac{\partial w}{\partial t} + u \frac{\partial w}{\partial x} + v \frac{\partial w}{\partial y} + w \frac{\partial w}{\partial z} + \frac{\partial P}{\partial z} - \frac{1}{Re} \left( \frac{\partial^2 w}{\partial x^2} + \frac{\partial^2 w}{\partial y^2} + \frac{\partial^2 w}{\partial z^2} \right) = 0$$

## Experimental data loss

$$\|\vec{V} - \vec{V}_{exp}\|^2 = 0$$

# Results



# Content

- Pure Data-driven Approaches
- Hybrid Approaches
  - Physics-Informed Diffusion Model
  - Physics-Informed Neural Network
  - Neural ODE
  - Lagrangian Network / Hamiltonian Network
  - Neural Operator
  - Fourier Neural Operator

# ODE Forms the Foundation of Simulation

- For example, positions are integrals of velocity

$$\frac{dx(t)}{dt} = v(t) \Leftrightarrow x(t) = x(0) + \int_0^t v(t) dt$$

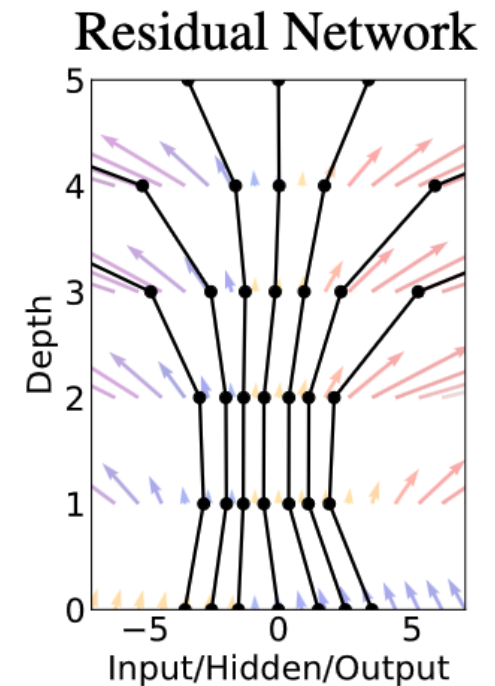

- Idea: specify an ordinary differential equation (ODE) by a neural network

$$\frac{dz(t)}{dt} = f(z(t), t, \theta) \Leftrightarrow z(t_1) = z(t_0) + \int_{t_0}^{t_1} f(z(t), t, \theta) dt$$

- How to learn the neural network?

# Parameterize Discretized ODE as Residual Network

- Parameterize discretized ODE:
$$\mathbf{z}_{t+1} = \mathbf{z}_t + f(\mathbf{h}_t, \theta_t)$$
- Learning objective: reconstruct data at every time step  $\mathbf{t}$ :  $\mathbf{L} = \sum_t \|\mathbf{z}_t - \hat{\mathbf{z}}_t\|^2$
- Problem:
  1. Requires a very deep network to handle long sequence
  2. Unable to simulate intermediate time steps

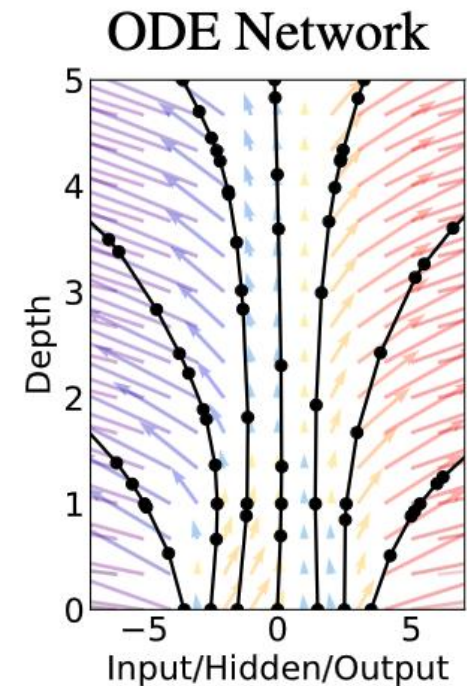



# Parameterize Continuous ODE as Neural Network

- Parameterize continuous ODE:

$$z(t_1) = z(t_0) + \int_{t_0}^{t_1} f(z(t), t, \theta) dt$$

- Learning objective: reconstruct data at subsampled time step  $t'$ :  $L = \sum_{t'} \|z_{t'} - \hat{z}_{t'}\|^2$
- Problem:
  - Gradients need to be back-propagated through time, which is memory inefficient



$$\begin{aligned} z(t + \Delta t) &= z(t) + f(z(t), t, \theta) \Delta t \\ z(t + \Delta t + \Delta t) &= z(t + \Delta t) + f(z(t + \Delta t), t + \Delta t, \theta) \Delta t \\ &= z(t) + f(z(t), t, \theta) \Delta t + f(z(t) + f(z(t), t, \theta) \Delta t, t + \Delta t, \theta) \Delta t \end{aligned}$$

# Backpropagate Gradients Through Time

- Simplify continuous ODE with infinitesimal interval  $\Delta t$

$$z(t + \Delta t) = z(t) + f(z(t), t, \theta)\Delta t$$

$$\begin{aligned} z(t + \Delta t + \Delta t) &= z(t + \Delta t) + f(z(t + \Delta t), t + \Delta t, \theta)\Delta t \\ &= z(t) + f(z(t), t, \theta)\Delta t + f(z(t) + f(z(t), t, \theta)\Delta t, t + \Delta t, \theta)\Delta t \end{aligned}$$

# Backpropagate Gradients Through Time

- Simplify continuous ODE with infinitesimal interval  $\Delta t$

$$z(t + \Delta t) = z(t) + f(z(t), t, \theta)\Delta t$$

$$\begin{aligned} z(t + \Delta t + \Delta t) &= z(t + \Delta t) + f(z(t + \Delta t), t + \Delta t, \theta)\Delta t \\ &= z(t) + f(z(t), t, \theta)\Delta t + f(z(t) + f(z(t), t, \theta)\Delta t, t + \Delta t, \theta)\Delta t \end{aligned}$$

- Gradients at  $t + \Delta t + \Delta t$ :

$$\frac{dL(t + \Delta t + \Delta t)}{d\theta} = \frac{dL(t + \Delta t + \Delta t)}{dz(t + \Delta t + \Delta t)} \frac{dz(t + \Delta t + \Delta t)}{d\theta}$$

Can we bypass gradient backpropagation through time?

---

# Neural Ordinary Differential Equations

---

**Ricky T. Q. Chen\***, **Yulia Rubanova\***, **Jesse Bettencourt\***, **David Duvenaud**  
University of Toronto, Vector Institute  
`{rtqichen, rubanova, jessebett, duvenaud}@cs.toronto.edu`


- Key idea: use the adjoint sensitivity method (Pontryagin et al., 1962) to bypass backpropagation through the ODE solver.
- The approach computes gradients by solving a second, augmented ODE backwards in time, and is applicable to all ODE solvers.

# Idea: Integrate ODE Neural Network into any ODE Solver

- Forward pass remains the same. Define a scalar loss function  $L(\cdot)$

$$L(\mathbf{z}(t_1)) = L\left(\mathbf{z}(t_0) + \int_{t_0}^{t_1} f(\mathbf{z}(t), t, \theta) dt\right) = L(\text{ODESolve}(\mathbf{z}(t_0), f, t_0, t_1, \theta))$$


- Calculate loss at sampled time steps



# Key Results

- First define adjoint:  $\mathbf{a}(t) = \partial \mathbf{L} / \partial \mathbf{z}(t)$  which determines how the gradient of the loss depends on the hidden state  $\mathbf{z}(t)$  at each instant
- Key result 1: the dynamics of  $\mathbf{a}(t)$  can be given by another ODE

$$\frac{d\mathbf{a}(t)}{dt} = -\mathbf{a}(t)^\top \frac{\partial f(\mathbf{z}(t), t, \theta)}{\partial \mathbf{z}}$$



This solver must run backwards, starting from the initial value of  $\partial \mathbf{L} / \partial \mathbf{z}(t_1)$ .

We can simply recompute  $\mathbf{z}(t)$  backwards in time together with the adjoint, starting from its final value  $\mathbf{z}(t_1)$

# Key Results

- First define adjoint:  $\mathbf{a}(t) = \partial L / \partial \mathbf{z}(t)$  which determines how the gradient of the loss depends on the hidden state  $\mathbf{z}(t)$  at each instant
- Key result 1: the dynamics of  $\mathbf{a}(t)$  can be given by another ODE

$$\frac{d\mathbf{a}(t)}{dt} = -\mathbf{a}(t)^\top \frac{\partial f(\mathbf{z}(t), t, \theta)}{\partial \mathbf{z}}$$

- Key result 2: computes the gradients with respect to parameter  $\theta$  requires evaluating a third integral, which depends on both  $\mathbf{z}(t)$  and  $\mathbf{a}(t)$ :

$$\frac{dL}{d\theta} = - \int_{t_1}^{t_0} \mathbf{a}(t)^\top \frac{\partial f(\mathbf{z}(t), t, \theta)}{\partial \theta} dt$$

# Proof of Result 1 $\frac{d\mathbf{a}(t)}{dt} = -\mathbf{a}(t)^\top \frac{\partial f(\mathbf{z}(t), t, \theta)}{\partial \mathbf{z}}$

- We first define:

$$\mathbf{a}(t) = \frac{dL}{d\mathbf{z}(t)} \quad \mathbf{z}(t + \varepsilon) = \int_t^{t+\varepsilon} f(\mathbf{z}(t), t, \theta) dt + \mathbf{z}(t) = T_\varepsilon(\mathbf{z}(t), t)$$

- From chain rule:

$$\frac{dL}{\partial \mathbf{z}(t)} = \frac{dL}{d\mathbf{z}(t + \varepsilon)} \frac{d\mathbf{z}(t + \varepsilon)}{d\mathbf{z}(t)} \quad \longrightarrow \quad \mathbf{a}(t) = \mathbf{a}(t + \varepsilon) \frac{\partial T_\varepsilon(\mathbf{z}(t), t)}{\partial \mathbf{z}(t)}$$

$$\frac{d\mathbf{a}(t)}{dt} = \lim_{\varepsilon \rightarrow 0^+} \frac{\mathbf{a}(t + \varepsilon) - \mathbf{a}(t)}{\varepsilon} \quad (39)$$

$$= \lim_{\varepsilon \rightarrow 0^+} \frac{\mathbf{a}(t + \varepsilon) - \mathbf{a}(t + \varepsilon) \frac{\partial}{\partial \mathbf{z}(t)} T_\varepsilon(\mathbf{z}(t))}{\varepsilon} \quad \text{(by Eq 38)} \quad (40)$$

$$= \lim_{\varepsilon \rightarrow 0^+} \frac{\mathbf{a}(t + \varepsilon) - \mathbf{a}(t + \varepsilon) \frac{\partial}{\partial \mathbf{z}(t)} (\mathbf{z}(t) + \varepsilon f(\mathbf{z}(t), t, \theta) + \mathcal{O}(\varepsilon^2))}{\varepsilon} \quad \text{(Taylor series around } \mathbf{z}(t)) \quad (41)$$

$$= \lim_{\varepsilon \rightarrow 0^+} \frac{\mathbf{a}(t + \varepsilon) - \mathbf{a}(t + \varepsilon) \left( I + \varepsilon \frac{\partial f(\mathbf{z}(t), t, \theta)}{\partial \mathbf{z}(t)} + \mathcal{O}(\varepsilon^2) \right)}{\varepsilon} \quad (42)$$

$$= \lim_{\varepsilon \rightarrow 0^+} \frac{-\varepsilon \mathbf{a}(t + \varepsilon) \frac{\partial f(\mathbf{z}(t), t, \theta)}{\partial \mathbf{z}(t)} + \mathcal{O}(\varepsilon^2)}{\varepsilon} \quad (43)$$

$$= \lim_{\varepsilon \rightarrow 0^+} -\mathbf{a}(t + \varepsilon) \frac{\partial f(\mathbf{z}(t), t, \theta)}{\partial \mathbf{z}(t)} + \mathcal{O}(\varepsilon) \quad (44)$$

$$= -\mathbf{a}(t) \frac{\partial f(\mathbf{z}(t), t, \theta)}{\partial \mathbf{z}(t)} \quad (45)$$

# Proof of Result 2

$$\frac{dL}{d\theta} = - \int_{t_1}^{t_0} \mathbf{a}(t)^\top \frac{\partial f(\mathbf{z}(t), t, \theta)}{\partial \theta} dt$$

- We have:  $\frac{\partial \theta(t)}{\partial t} = \mathbf{0}$      $\frac{dt(t)}{dt} = 1$

- We define augmented ODE:

$$\frac{d}{dt} \begin{bmatrix} \mathbf{z} \\ \theta \\ t \end{bmatrix} (t) = f_{aug}([\mathbf{z}, \theta, t]) := \begin{bmatrix} f([\mathbf{z}, \theta, t]) \\ \mathbf{0} \\ 1 \end{bmatrix}, \quad \mathbf{a}_{aug} := \begin{bmatrix} \mathbf{a} \\ \mathbf{a}_\theta \\ \mathbf{a}_t \end{bmatrix}, \quad \mathbf{a}_\theta(t) := \frac{dL}{d\theta(t)}, \quad \mathbf{a}_t(t) := \frac{dL}{dt(t)}$$

- We have

$$\frac{\partial f_{aug}}{\partial [\mathbf{z}, \theta, t]} = \begin{bmatrix} \frac{\partial f}{\partial \mathbf{z}} & \frac{\partial f}{\partial \theta} & \frac{\partial f}{\partial t} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix} (t)$$

# Proof of Result 2

$$\frac{dL}{d\theta} = - \int_{t_1}^{t_0} \mathbf{a}(t)^\top \frac{\partial f(\mathbf{z}(t), t, \theta)}{\partial \theta} dt$$

- From key result 1, we have  $\frac{d\mathbf{a}(t)}{dt} = -\mathbf{a}(t)^\top \frac{\partial f(\mathbf{z}(t), t, \theta)}{\partial \mathbf{z}}$
- Substitute it to  $\mathbf{a}_{aug}$

$$\frac{d\mathbf{a}_{aug}(t)}{dt} = - [\mathbf{a}(t) \quad \mathbf{a}_\theta(t) \quad \mathbf{a}_t(t)] \frac{\partial f_{aug}}{\partial [\mathbf{z}, \theta, t]}(t) = - \left[ \mathbf{a} \frac{\partial f}{\partial \mathbf{z}} \quad \mathbf{a} \frac{\partial f}{\partial \theta} \quad \mathbf{a} \frac{\partial f}{\partial t} \right] (t)$$

- As a result, we have:

$$\frac{dL}{d\theta} = \mathbf{a}_\theta(t_0) = - \int_{t_N}^{t_0} \mathbf{a}(t) \frac{\partial f(\mathbf{z}(t), t, \theta)}{\partial \theta} dt$$

# Overall Algorithm

---

**Algorithm 1** Reverse-mode derivative of an ODE initial value problem

---

**Input:** dynamics parameters  $\theta$ , start time  $t_0$ , stop time  $t_1$ , final state  $\mathbf{z}(t_1)$ , loss gradient  $\frac{\partial L}{\partial \mathbf{z}(t_1)}$   
 $s_0 = [\mathbf{z}(t_1), \frac{\partial L}{\partial \mathbf{z}(t_1)}, \mathbf{0}_{|\theta|}]$  ▷ Define initial augmented state  
**def** `aug_dynamics`(`[\mathbf{z}(t), \mathbf{a}(t), \cdot]`, `t`, `\theta`): ▷ Define dynamics on augmented state  
    **return** `[f(\mathbf{z}(t), t, \theta), -\mathbf{a}(t)^\top \frac{\partial f}{\partial \mathbf{z}}, -\mathbf{a}(t)^\top \frac{\partial f}{\partial \theta}]` ▷ Compute vector-Jacobian products  
`[\mathbf{z}(t_0), \frac{\partial L}{\partial \mathbf{z}(t_0)}, \frac{\partial L}{\partial \theta}] = \text{ODESolve}(s_0, \text{aug\_dynamics}, t_1, t_0, \theta)` ▷ Solve reverse-time ODE  
**return** `\frac{\partial L}{\partial \mathbf{z}(t_0)}, \frac{\partial L}{\partial \theta}` ▷ Return gradients

---


1. Obtain the dynamics of  $\mathbf{a}(t)$  by reverse ODE

$$\frac{d\mathbf{a}(t)}{dt} = -\mathbf{a}(t)^\top \frac{\partial f(\mathbf{z}(t), t, \theta)}{\partial \mathbf{z}}$$




2. Compute the gradients with respect to parameter  $\theta$  with a third integral

$$\frac{dL}{d\theta} = - \int_{t_1}^{t_0} \mathbf{a}(t)^\top \frac{\partial f(\mathbf{z}(t), t, \theta)}{\partial \theta} dt$$



# Results



# Results



(a) Recurrent Neural Network



(b) Latent Neural Ordinary Differential Equation

# Content

- Pure Data-driven Approaches
- Hybrid Approaches
  - Physics-Informed Diffusion Model
  - Physics-Informed Neural Network
  - Neural ODE
  - Lagrangian Network / Hamiltonian Network
  - Neural Operator
  - Fourier Neural Operator

# Lagrangian Mechanics

- The Lagrangian:


Potential energy

$$L = T - V$$

Kinetic energy

- Euler-Lagrange Equation:

$$\frac{d}{dt} \frac{\partial L}{\partial \dot{\mathbf{q}}} = \frac{\partial L}{\partial \mathbf{q}} \Rightarrow \frac{d}{dt} \frac{\partial T}{\partial \dot{\mathbf{q}}} = - \frac{\partial V}{\partial \mathbf{q}}$$



Joseph-Louis Lagrange


# Hamiltonian Mechanics

- We can derive Hamiltonian mechanics from Lagrangian:

$$p_i = \frac{\delta L}{\delta \dot{q}_i}; \quad H = \sum_i \dot{q}_i p_i - L$$

- Hamiltonian Equation:

$$\frac{dq}{dt} = \frac{\partial H}{\partial p}; \quad \frac{dp}{dt} = -\frac{\partial H}{\partial q}$$



William Rowan Hamilton

Can We Directly Bake Lagrangian / Hamiltonian  
Mechanics into Neural Networks?


# Hamiltonian Mechanics

- We can derive Hamiltonian mechanics from Lagrangian:

$$p_i = \frac{\delta L}{\delta \dot{q}_i}; \quad H = \sum_i \dot{q}_i p_i - L$$

- Hamiltonian Equation:

$$\frac{dq}{dt} = \frac{\partial H}{\partial p}; \quad \frac{dp}{dt} = -\frac{\partial H}{\partial q}$$

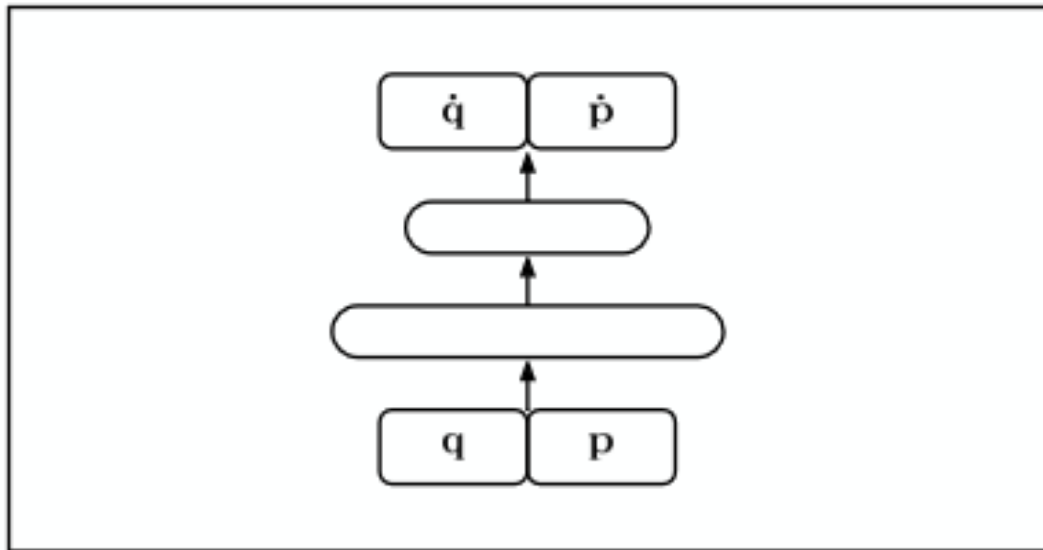


William Rowan Hamilton

# Output Motions vs. Output Energy to Derive Motions


- Hamiltonian Equation:  $\frac{dq}{dt} = \frac{\partial H}{\partial p}$ ;  $\frac{dp}{dt} = -\frac{\partial H}{\partial q}$

Baseline NN



Output motions

Hamiltonian NN



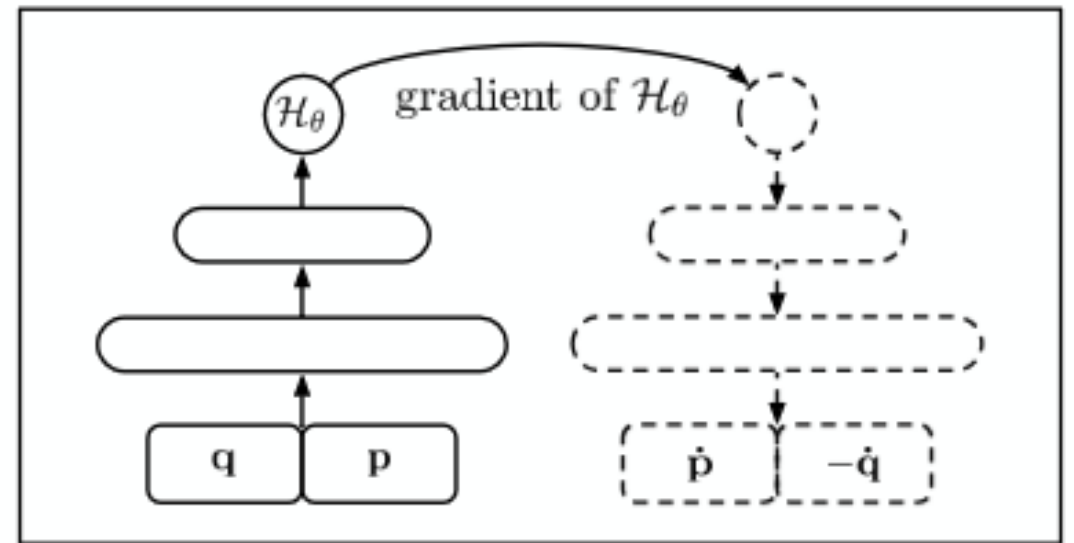
Output energy and derive motions with autograd

# Output Motions vs. Output Energy to Derive Motions

- Hamiltonian Equation:  $\frac{dq}{dt} = \frac{\partial H}{\partial p}$ ;  $\frac{dp}{dt} = -\frac{\partial H}{\partial q}$

$$\mathcal{L}_{HNN} = \left\| \frac{\partial \mathcal{H}_\theta}{\partial \mathbf{p}} - \frac{\partial \mathbf{q}}{\partial t} \right\|_2 + \left\| \frac{\partial \mathcal{H}_\theta}{\partial \mathbf{q}} + \frac{\partial \mathbf{p}}{\partial t} \right\|_2$$

Hamiltonian NN



Output energy and derive motions with autograd

# Results

Ground truth

Baseline NN

Hamiltonian NN




# Task 1: Ideal Mass-Spring System


Hamiltonian

$$\mathcal{H} = \frac{1}{2}kq^2 + \frac{p^2}{2m}$$

Ideal mass-spring system




Noisy observations



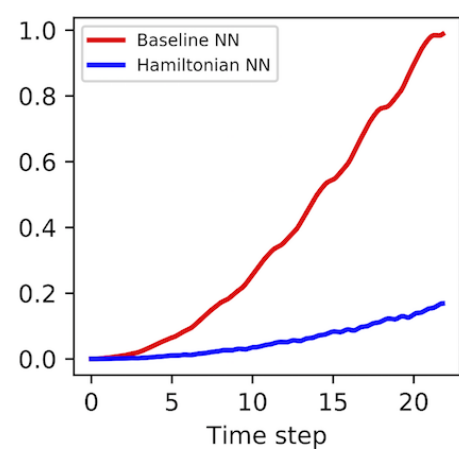
Training settings

- 25 trajectories of 30 observations each
- Adam with  $10^{-3}$  learn rate
- 3 layers of 200 units
- 2000 gradient steps
- $\tanh$  activations


Predictions




MSE between coordinates



Total HNN-conserved quantity



Total energy




Energy is conserved!

# Task 3: Real Pendulum


Hamiltonian

$$\mathcal{H} = 2mgl(1 - \cos q) + \frac{l^2 p^2}{2m}$$

Real pendulum




Noisy observations




Training settings

- 1 trajectory of 555 observations
- Adam with  $10^{-3}$  learn rate
- 3 layers of 200 units
- 2000 gradient steps
- `tanh` activations


Predictions



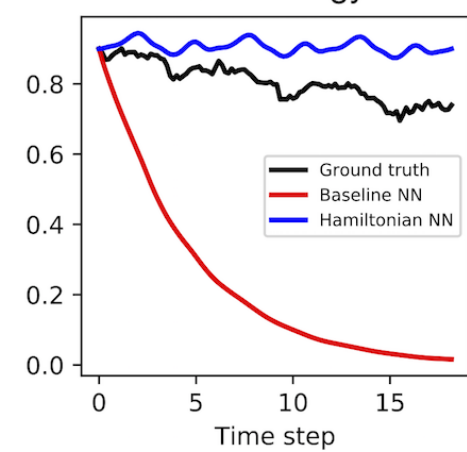
MSE between coordinates



Total HNN-conserved quantity



Total energy



Energy is conserved!

# Lagrangian Mechanics

- The Lagrangian:


Potential energy

$$L = T - V$$

Kinetic energy

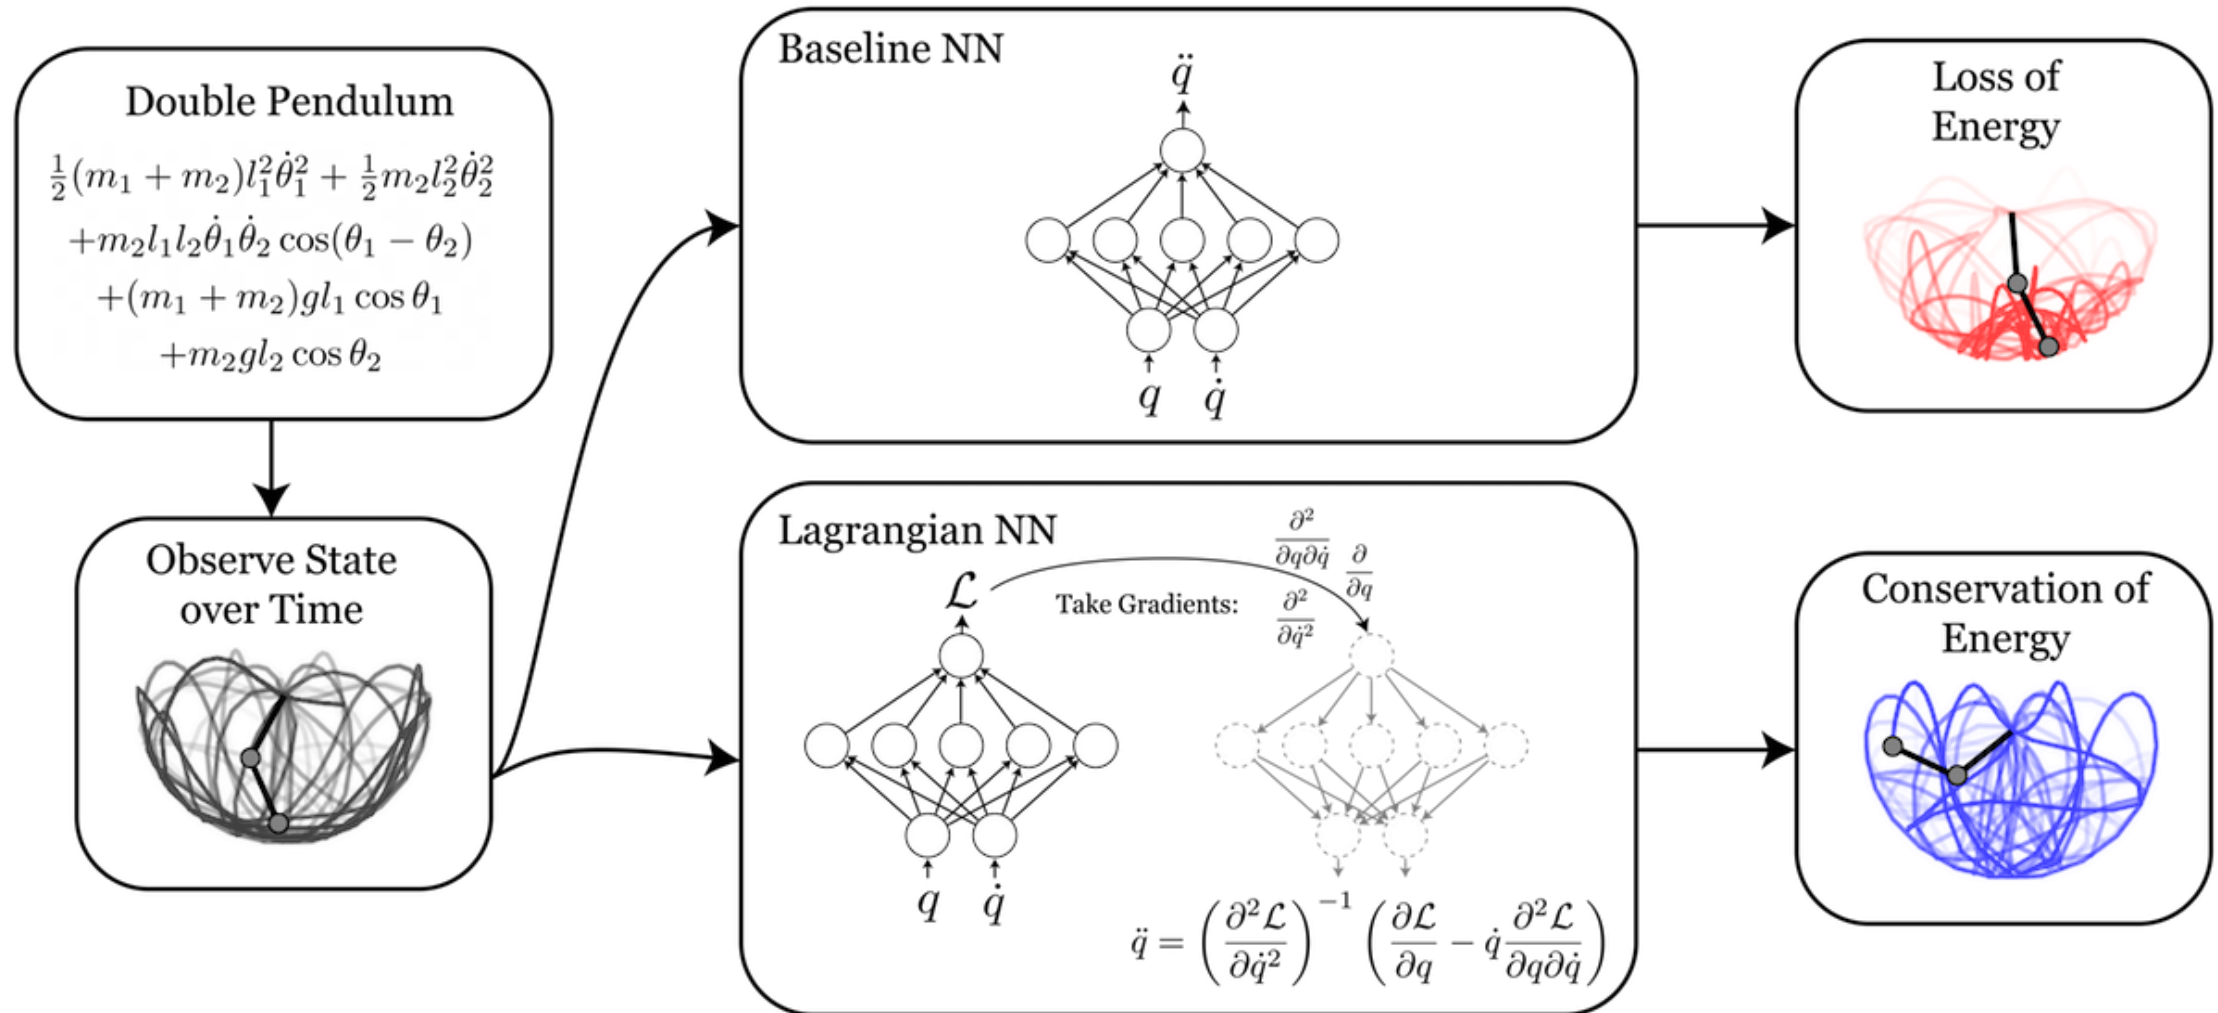
- Euler-Lagrange Equation:

$$\frac{d}{dt} \frac{\partial L}{\partial \dot{\mathbf{q}}} = \frac{\partial L}{\partial \mathbf{q}} \Rightarrow \frac{d}{dt} \frac{\partial T}{\partial \dot{\mathbf{q}}} = - \frac{\partial V}{\partial \mathbf{q}}$$



Joseph-Louis Lagrange

# Lagrangian Neural Network Follows the Same Idea



# The Derivation of Motions from Lagrangian

$$\frac{d}{dt} \frac{\partial \mathcal{L}}{\partial \dot{q}_j} = \frac{\partial \mathcal{L}}{\partial q_j} \quad \text{the Euler-Lagrange equation} \quad (5)$$


$$\frac{d}{dt} \nabla_{\dot{q}} \mathcal{L} = \nabla_q \mathcal{L} \quad \text{switch to vector notation} \quad (6)$$

$$(\nabla_{\dot{q}} \nabla_{\dot{q}}^\top \mathcal{L}) \ddot{q} + (\nabla_q \nabla_{\dot{q}}^\top \mathcal{L}) \dot{q} = \nabla_q \mathcal{L} \quad \text{expand time derivative } \frac{d}{dt} \quad (7)$$

$$\ddot{q} = (\nabla_{\dot{q}} \nabla_{\dot{q}}^\top \mathcal{L})^{-1} [\nabla_q \mathcal{L} - (\nabla_q \nabla_{\dot{q}}^\top \mathcal{L}) \dot{q}] \quad \text{matrix inverse to solve for } \ddot{q} \quad (8)$$

# Results

Baseline NN



Lagrangian NN




# Content

- Pure Data-driven Approaches
- Hybrid Approaches
  - Physics-Informed Diffusion Model
  - Physics-Informed Neural Network
  - Neural ODE
  - Lagrangian Network / Hamiltonian Network
  - Neural Operator
  - Fourier Neural Operator

# Function Learning

- Function learning: given an input point  $\mathbf{y}$ , learn to predict  $\mathbf{f}(\mathbf{y})$
- For example, the function  $\mathbf{f}(\cdot)$  describes the fluid motion




The model only learns to simulate one function



Video source: Genesis simulator

# The Idea of Functional

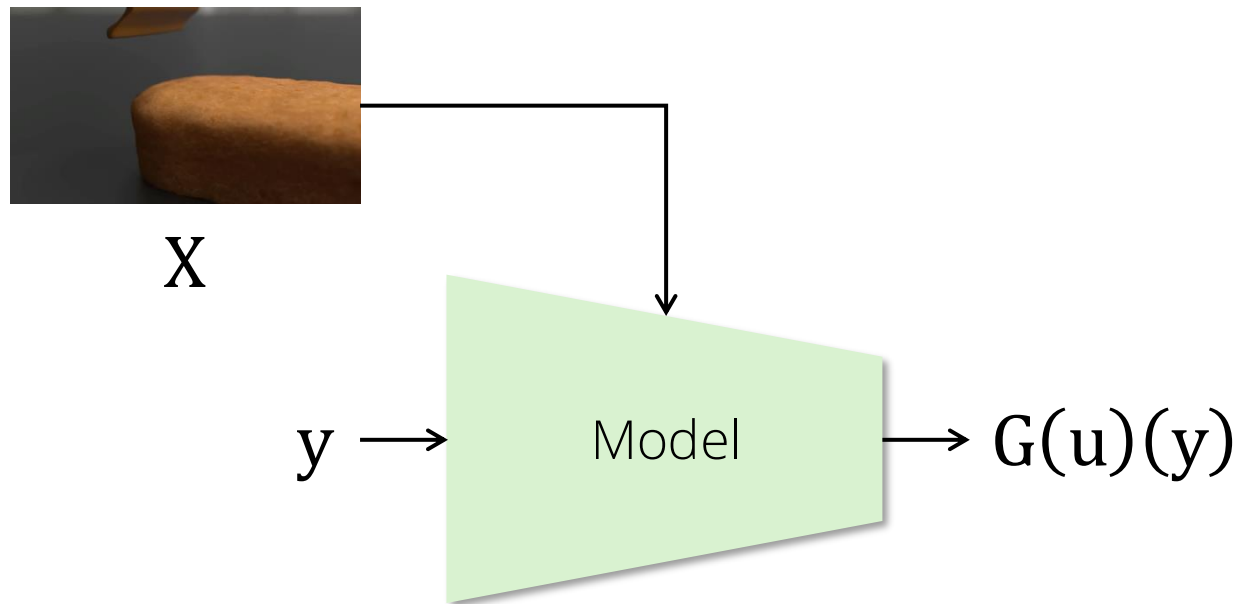
- Let  $\mathbf{G}$  denote an operator taking an input function  $\mathbf{u}$ , then  $\mathbf{G}(\mathbf{u})$  is the corresponding output function and  $\mathbf{G}(\mathbf{u})(\mathbf{y})$  evaluates the function at point  $\mathbf{y}$
- For example, the function  $\mathbf{G}(\text{fluid})$ ,  $\mathbf{G}(\text{sponge})$ ,  $\mathbf{G}(\text{fruit})$  describes the fluid, sponge and fruit motion



Video source: Genesis simulator

# Extend Function Learning to Operator Learning

- Operator learning: given an input function  $\mathbf{u}$  and an input point  $\mathbf{y}$ , learn to predict  $\mathbf{G}(\mathbf{u})(\mathbf{y})$
- The input function  $\mathbf{u}$  is obtained from a sparse set of observations  $\mathbf{X}$ .
- For example,  $\mathbf{X}$  may be the video of the fluid motion and  $\mathbf{y}$  corresponds to an internal point of the fluid.



## Why Operator Learning is powerful?


- Share structures among different dynamics
- Decouple conditional inputs from outputs
  - Predict high-res outputs while conditioning on low-res inputs
  - Conditions inputs from different modalities (e.g. lidar, radar, video ...)

# DeepONet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators

Lu Lu<sup>1</sup>, Pengzhan Jin<sup>2</sup>, and George Em Karniadakis<sup>1</sup>

<sup>1</sup>Division of Applied Mathematics, Brown University, Providence, RI 02912, USA


<sup>2</sup>LSEC, ICMSEC, Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Beijing 100190, China




# A Tensor-Product Factorization

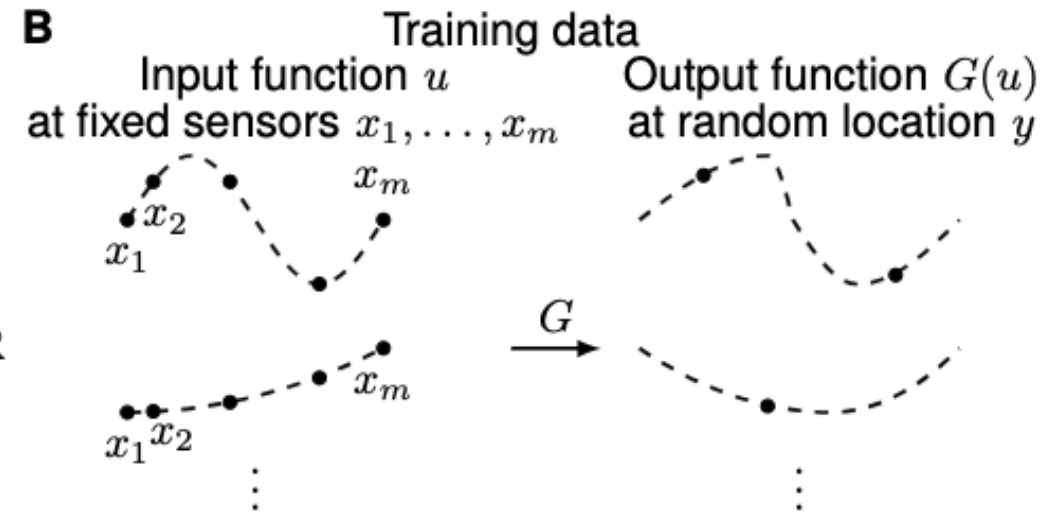

- Idea: predict  $G(\mathbf{u})(\mathbf{y})$  with the factorization of  $\mathbf{h}(\mathbf{X}) \otimes \mathbf{h}'(\mathbf{y})$  since tensor product  $\mathbf{V} \otimes \mathbf{W}$  forms the most general vector space that receives a bilinear map from  $\mathbf{V} \times \mathbf{W}$

$$\left| G(\mathbf{u})(\mathbf{y}) - \underbrace{\sum_{k=1}^p \sum_{i=1}^n c_i^k \sigma \left( \sum_{j=1}^m \xi_{ij}^k u(x_j) + \theta_i^k \right)}_{\text{branch}} \underbrace{\sigma(w_k \cdot \mathbf{y} + \zeta_k)}_{\text{trunk}} \right| < \epsilon$$



# Results





- Pros: query the output at any point
- Cons: the model doesn't generalize to different resolutions of observations of input functions

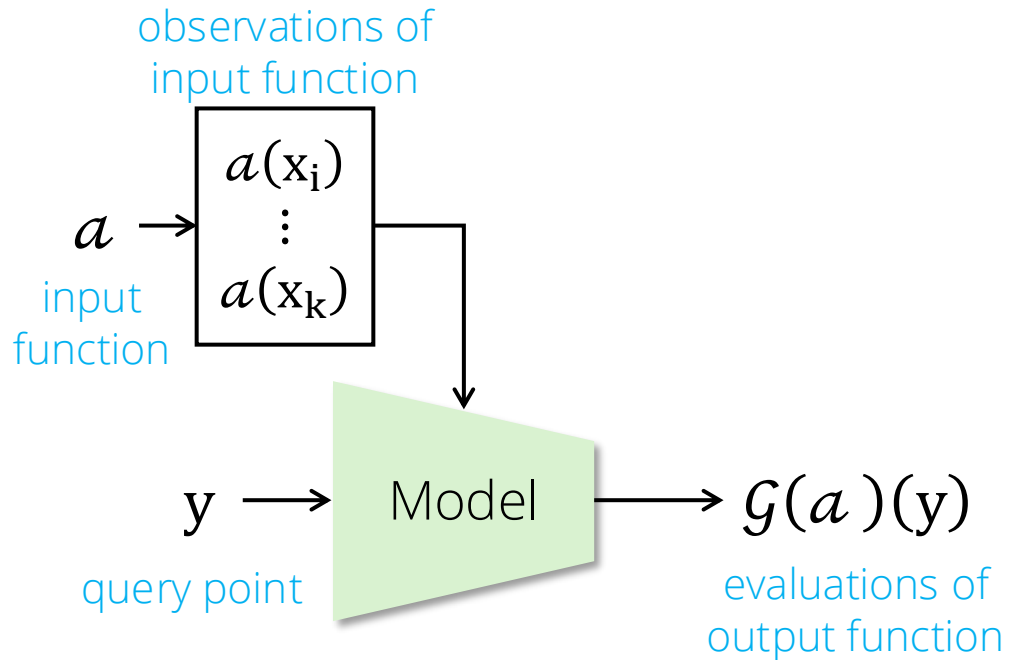
Can we build a neural operator that is Discretization Invariant?

# Content

- Pure Data-driven Approaches
- Hybrid Approaches
  - Physics-Informed Diffusion Model
  - Physics-Informed Neural Network
  - Neural ODE
  - Lagrangian Network / Hamiltonian Network
  - Neural Operator
  - Fourier Neural Operator


# Conceptual Comparison

## DeepONet



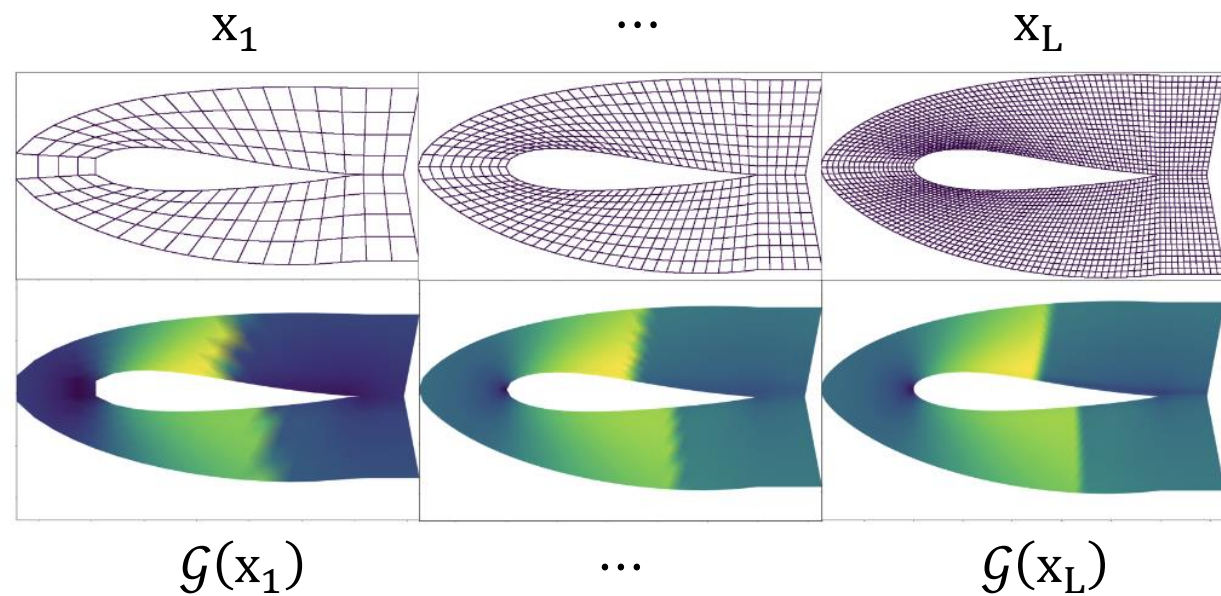
## Neural Operator:

Convert an input function  $a$  to an output function  $\mathcal{G}(a)$



# Neural Operator in Reality


Sampled points at different levels of discretization



Evaluations of output functions at different levels of discretization

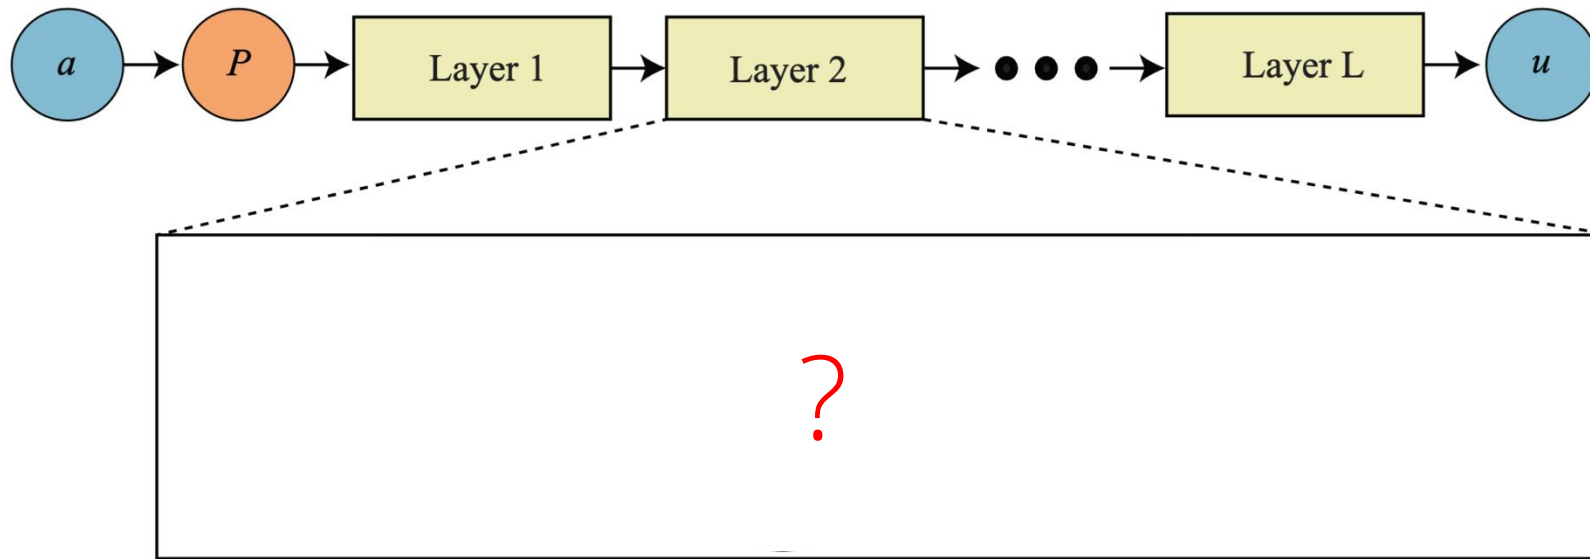
Neural Operator:

Convert an input function  $a$  to an output function  $\mathcal{G}(a)$




Can we train on one level of discretization while generalizing to different levels of discretization?


# Map One Function to Another while Maintaining Discretization Invariant




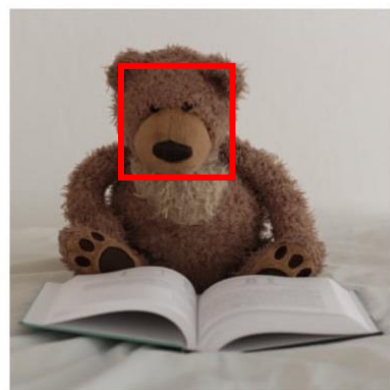
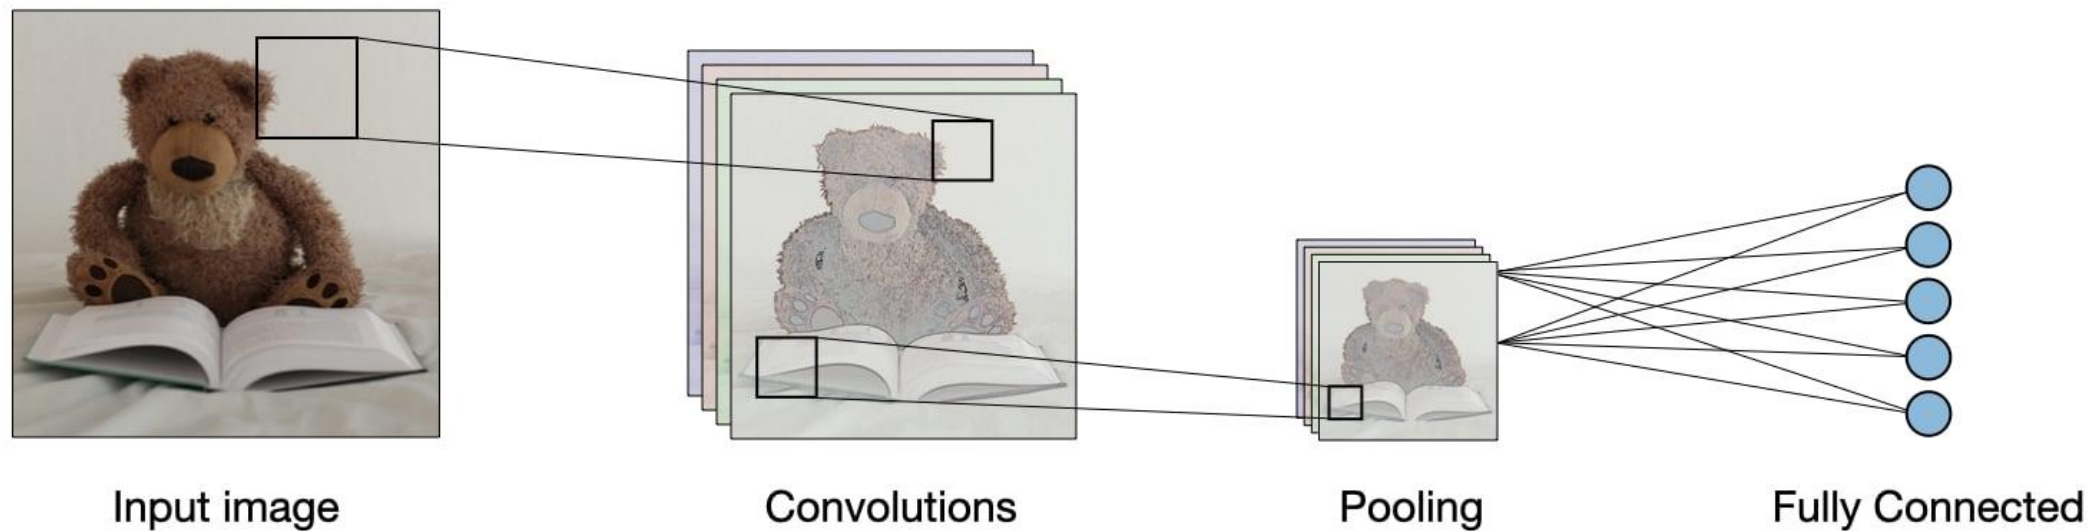
Train



Test



# Methods Function on Grid is not Discretization Invariant



# Key Idea: Learn a Point-wise Kernel

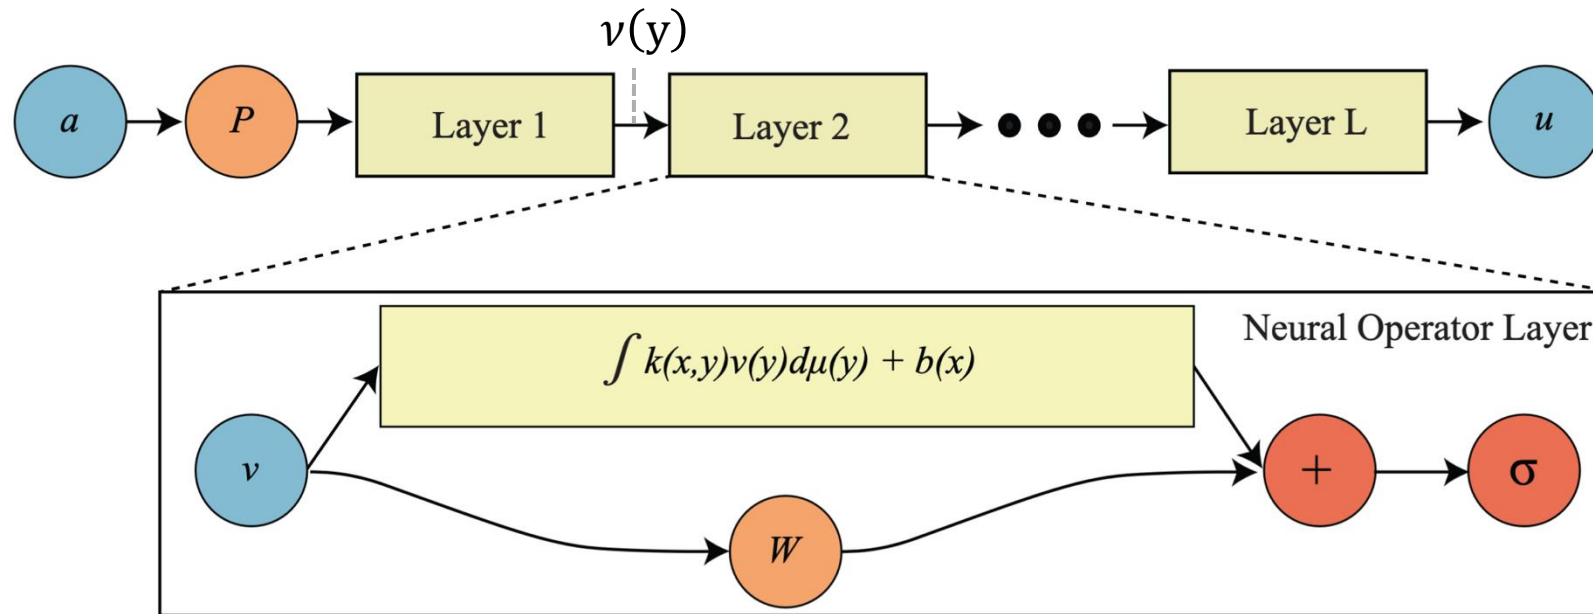
$$\int \kappa_{\phi}(\mathbf{x}, \mathbf{y}) v(\mathbf{y}) d\mathbf{y}$$

$\kappa_{\phi}(\mathbf{x}, \mathbf{y})$  denotes a kernel function that takes as input a pair of points  $(\mathbf{x}, \mathbf{y})$  and outputs a scalar

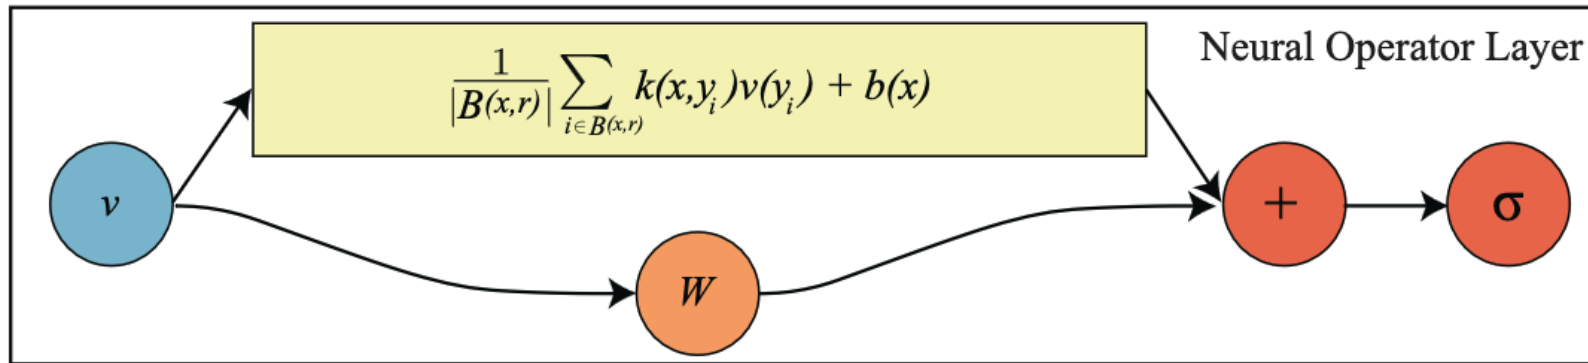
How to ensure discretization invariance?  
If  $\kappa_{\phi}$  is a continuous kernel, optimizing it with subsampled  $\mathbf{x}_l$  fits  $\kappa_{\phi}$  on the entire "continuous" domain of  $\mathbf{x}$

$\kappa_{\phi}(\mathbf{x}, \mathbf{y}) v(\mathbf{y})$  aggregates features from every point  $\mathbf{y}$  for the query point  $\mathbf{x}$

# The Diagram of Neural Operator




# Choice 1: A Shallow Neural Layer

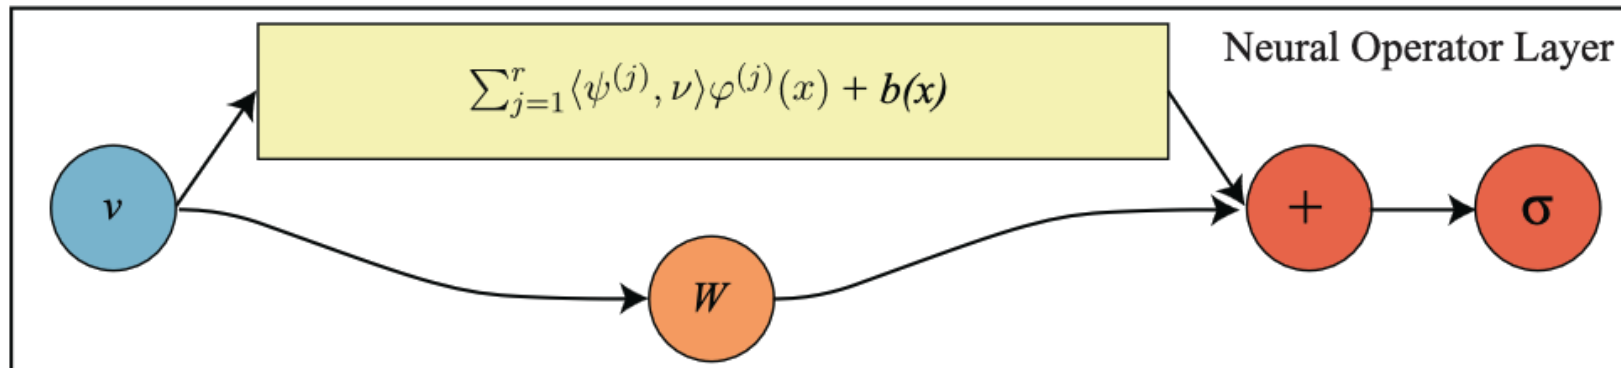


- Let  $\mathbf{n}$  be the input channel dimension,  $\mathbf{m}$  be the output channel dimensions, and  $\mathbf{d}$  be the coordinate dimension of point  $\mathbf{x}, \mathbf{y}$
- Define  $\kappa_\phi(\mathbf{x}, \mathbf{y})$  as a kernel matrix  $\mathbf{K} \in \mathbb{R}^{\mathbf{m}\mathbf{d} \times \mathbf{n}\mathbf{d}}$

$$\kappa_\phi(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{K}_{ij} \in \mathbb{R}^{\mathbf{m} \times \mathbf{n}}$$



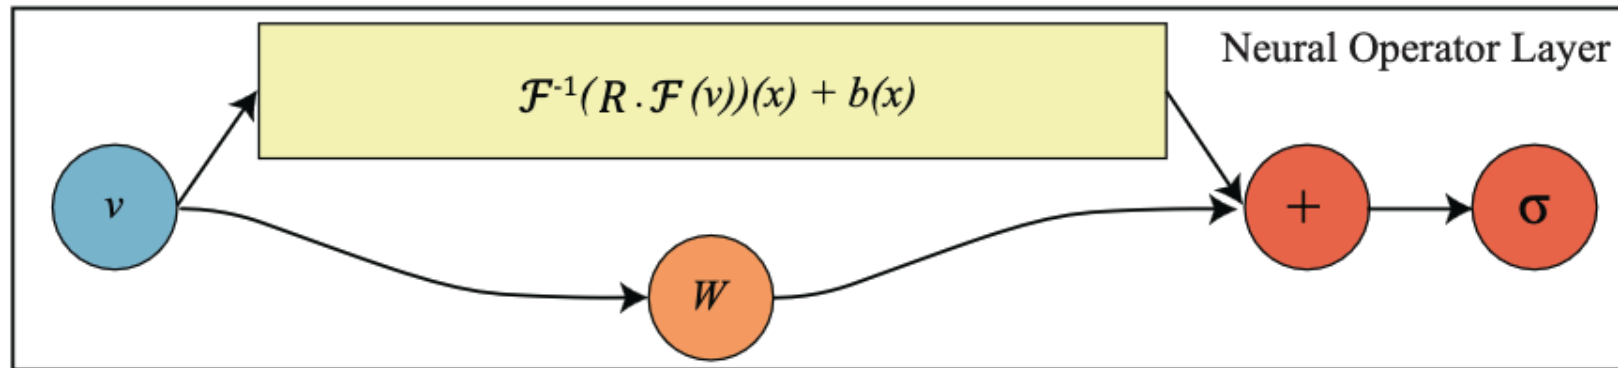
# Choice 2: Tensor Product Factorization



- Let  $\varphi^{(j)}: \mathbb{R}^d \rightarrow \mathbb{R}^m$  and  $\psi^{(j)}: \mathbb{R}^d \rightarrow \mathbb{R}^n$
- Define  $\kappa_\phi(\mathbf{x}, \mathbf{y})$  as a summation of tensor products

$$\kappa_\phi(\mathbf{x}, \mathbf{y}) = \sum_j \varphi^{(j)}(\mathbf{x}) \otimes \psi^{(j)}(\mathbf{y})$$

# Choice 3: Spectral Factorization



- Define  $\kappa_\phi(\mathbf{x}, \mathbf{y})$  in the Fourier domain  $R_\phi$

# Choice 3: Spectral Factorization

- Let's first assume  $\kappa_\phi(\mathbf{x}, \mathbf{y})$  is translation invariant:

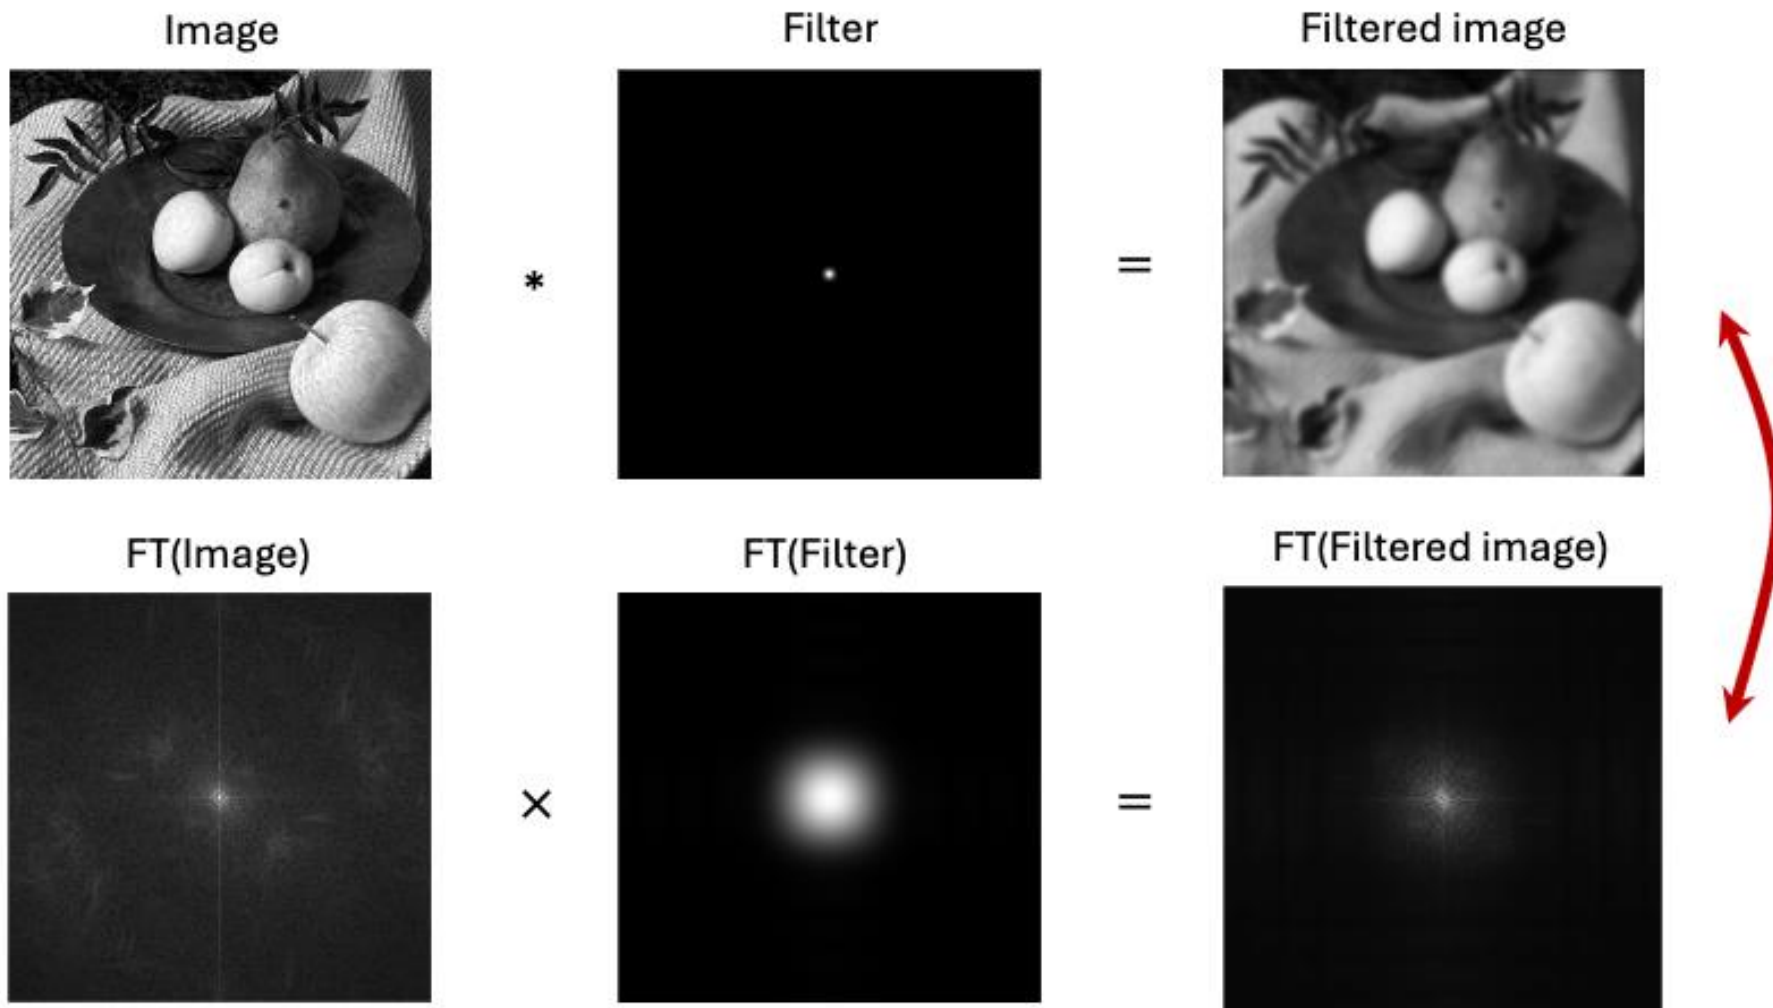
$$\kappa_\phi(\mathbf{x}, \mathbf{y}) = \kappa_\phi(\mathbf{x} - \mathbf{y})$$

- Therefore, the kernel operation is equivalent to convolutional operation

$$u(\mathbf{x}) = \int \kappa_\phi(\mathbf{x}, \mathbf{y})v(\mathbf{y})d\mathbf{y} = \int \kappa_\phi(\mathbf{x} - \mathbf{y})v(\mathbf{y})d\mathbf{y} = (\kappa * v)(\mathbf{x})$$

- By convolution theorem, the Fourier transform of a convolution of two functions is the product of their Fourier transforms

# 2D convolution theorem example



# Fourier analysis

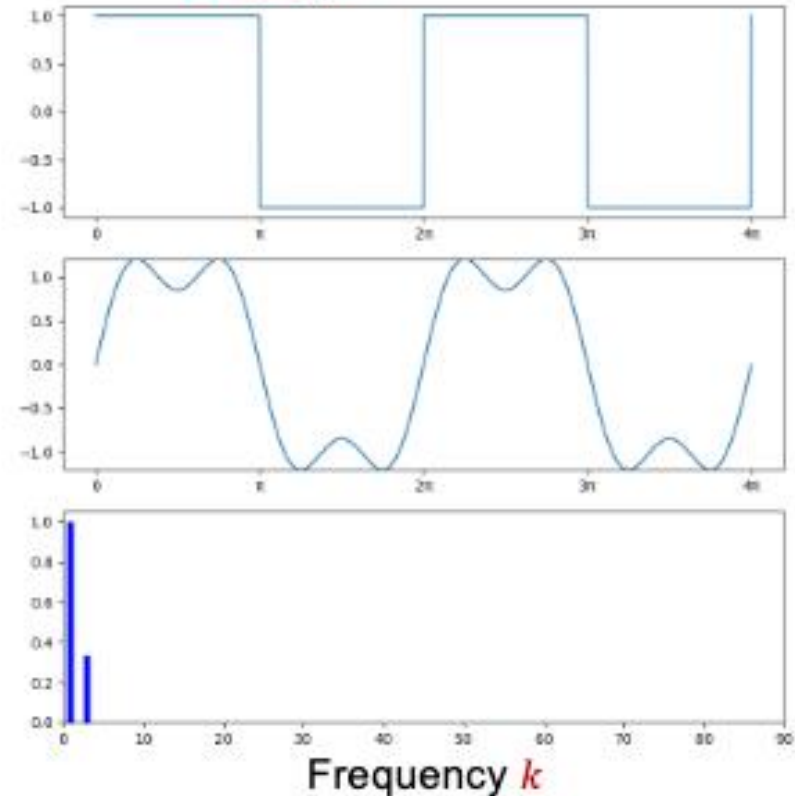
- Any “reasonable” univariate function can be expressed as a weighted sum of sinusoids of different frequencies (1807)



Jean-Baptiste Joseph Fourier (1768-1830)

Example: series for a square wave

$$\sum_{k=1,3,5,\dots}^{\infty} \frac{1}{k} \sin(kt)$$



# 1D Fourier transform

---

- Let's define basis functions parameterized by  $u \in (-\infty, \infty)$ :

$$\psi_u(t) = e^{i2\pi ut}$$

- Inner product for complex functions is given by:

$$\langle f, g \rangle = \int_{-\infty}^{\infty} f(t)g^*(t)dt$$

- Orthonormality:

$$\langle \psi_{u_1}, \psi_{u_2} \rangle = \begin{cases} 1 & \text{if } u_1 = u_2 \\ 0 & \text{otherwise} \end{cases}$$

# 1D Fourier transform

---

- Given a signal  $f(t)$ , we want to represent it as a weighted combination of the basis functions  $\psi_u(t) = e^{i2\pi ut}$  with weights  $F(u)$ :

$$f(t) = \int_{-\infty}^{\infty} F(u) e^{i2\pi ut} du$$

- Each weight  $F(u)$  is given by the inner product of  $f$  and  $\psi_u$ :

$$F(u) = \langle f, \psi_u \rangle = \int_{-\infty}^{\infty} f(t) e^{-i2\pi ut} dt$$

# 1D Fourier transform

---

- Forward transform:  $f(t) \xrightarrow{\mathcal{F}} F(u)$


$$F(u) = \int_{-\infty}^{\infty} f(t) e^{-i2\pi ut} dt$$

- Inverse transform:  $F(u) \xrightarrow{\mathcal{F}^{-1}} f(t)$

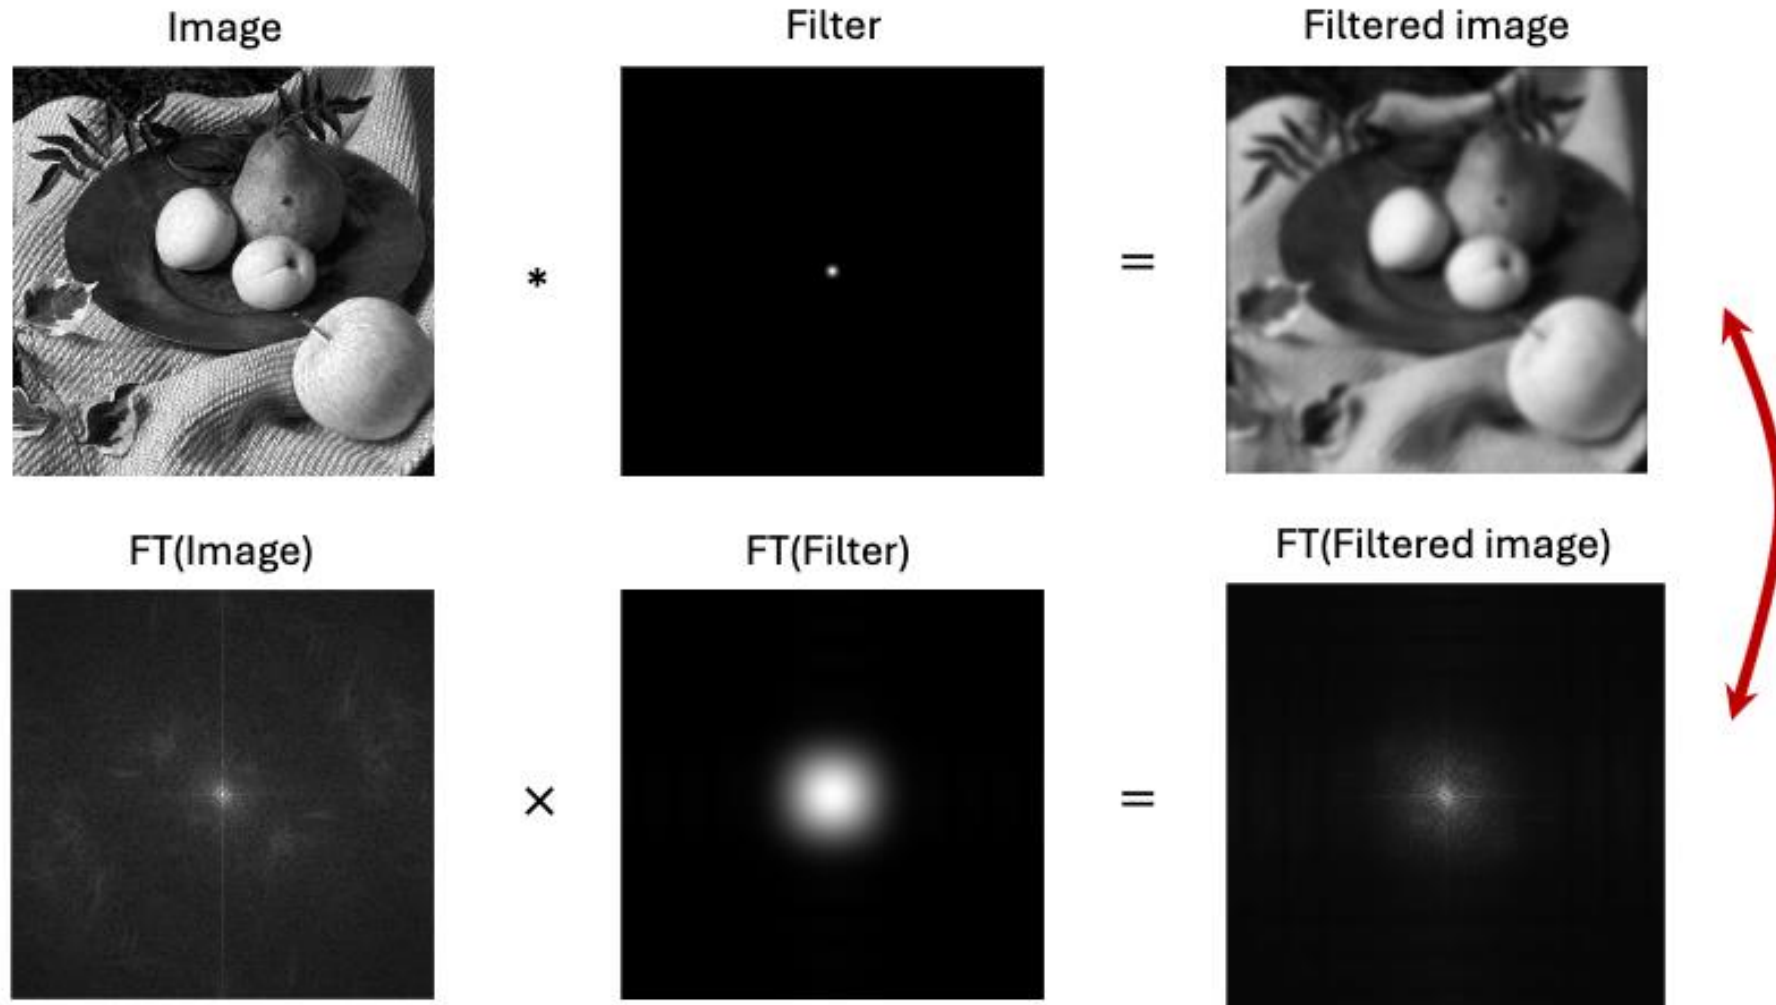
$$f(t) = \int_{-\infty}^{\infty} F(u) e^{i2\pi ut} du$$

- Duality: if  $f(t) \xrightarrow{\mathcal{F}} F(u)$ , then  $F(t) \xrightarrow{\mathcal{F}} f(-u)$

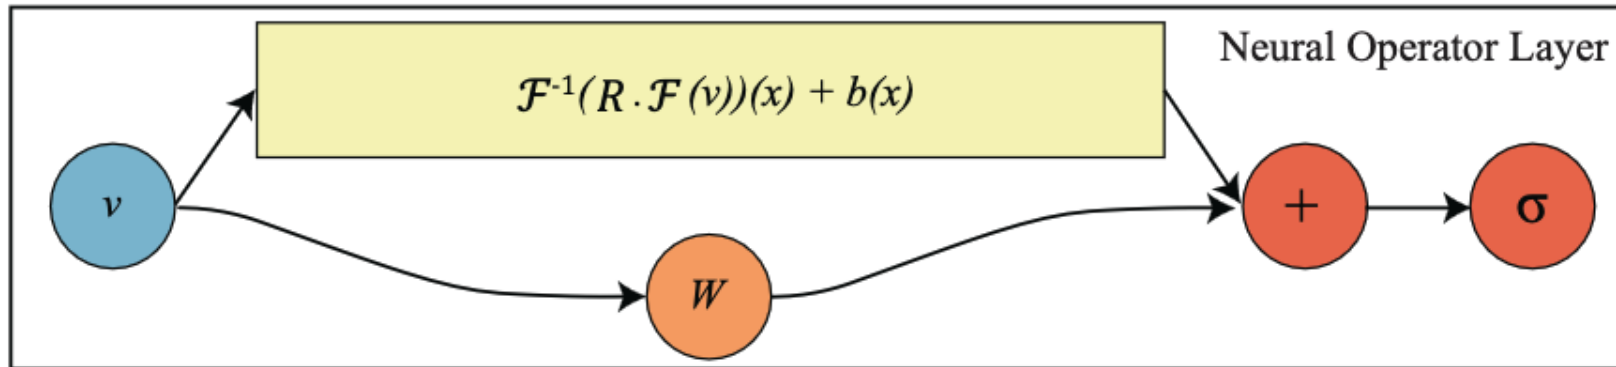
- Thus, we can talk about *Fourier transform pairs*  $f(t) \leftrightarrow F(u)$



# 2D convolution theorem example



# Choice 3: Spectral Factorization



- Define  $\kappa_{\phi}(\mathbf{x}, \mathbf{y})$  in the Fourier domain  $R_{\phi}$
- Therefore

$$\mathbf{u}(\mathbf{x}) = (\kappa * \mathbf{v})(\mathbf{x}) = \mathcal{F}^{-1} \left( R_{\phi} \cdot \mathcal{F}(\mathbf{v}) \right) (\mathbf{x})$$

Why spectral factorization? We can make the computation very efficient by truncating to few lowest-frequency modes

# 1D Fourier transform

- Forward transform:  $f(t) \xrightarrow{\mathcal{F}} F(u)$

$$F(u) = \int_{-\infty}^{\infty} f(t) e^{-i2\pi ut} dt$$

$-k_{\max} \leq u \leq k_{\max}$


- Inverse transform:  $F(u) \xrightarrow{\mathcal{F}^{-1}} f(t)$

$$f(t) = \int_{-\infty}^{\infty} F(u) e^{i2\pi ut} du$$

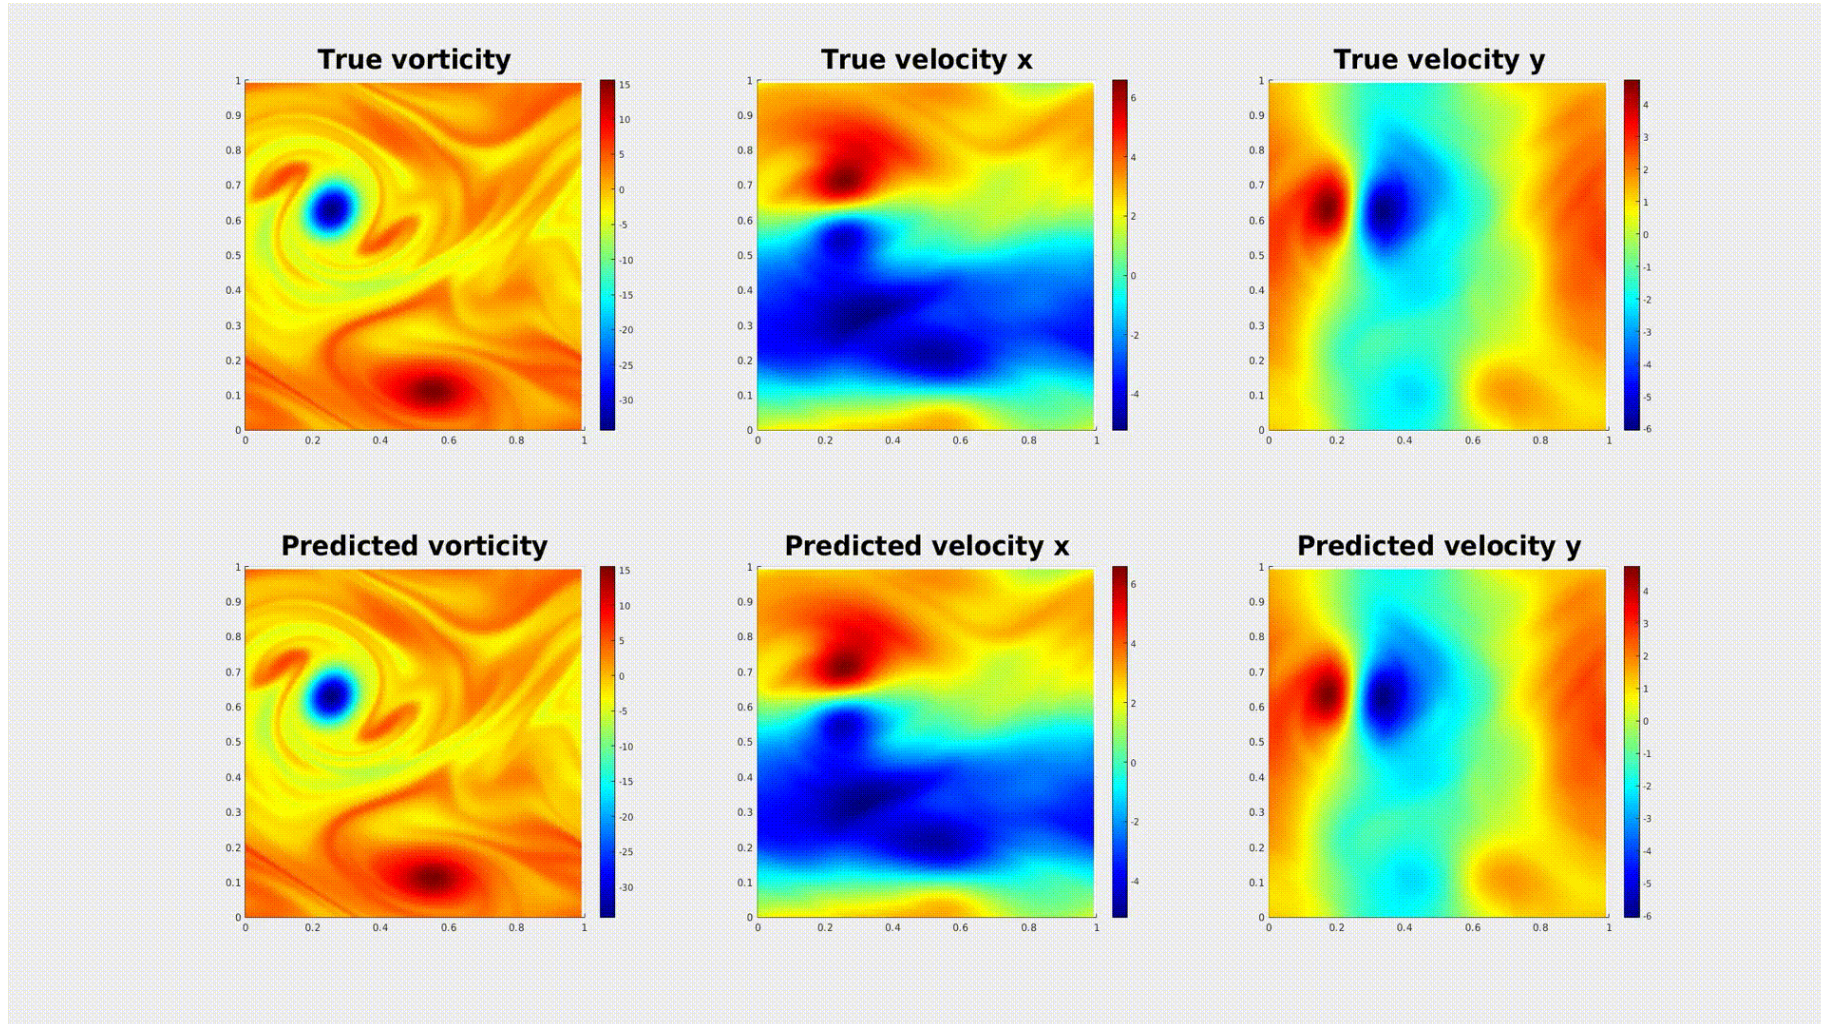
$k_{\max}$   
 $-k_{\max}$

- Duality: if  $f(t) \xrightarrow{\mathcal{F}} F(u)$ , then  $F(t) \xrightarrow{\mathcal{F}} f(-u)$

- Thus, we can talk about *Fourier transform pairs*  $f(t) \leftrightarrow F(u)$




# Results




# SFNO predictions on spherical shallow water equations

Train


Input x (32, 64)



Ground-truth y




SFNO prediction




Test


Input x (64, 128)



Ground-truth y




SFNO prediction




# Train

FNO predictions on 16x16 Darcy-Flow data



# Test

Zero-shot super-resolution: 16x16  $\rightarrow$  32x32



# Quick Summary

- To model the physical world, we've cover
  - 3D/4D modeling
  - Generative modeling
  - Rigid-body motion
  - Particle dynamics and mass-spring systems
  - Lagrangian dynamics and Incremental Potential Contact
  - Continuum dynamics and Material Point Method
  - Learning-based dynamics

# What We Will Cover the Next Week

- Embodied perception
  - Perceptive Action Decision
  - Multi-sensory Perception
  - Visual language model for embodied perception