

Partial Differential Equations and Neural Operators

Zongyi Li May 2022

CS159: Representation Learning for Science Prof. Yisong Yue

Overview

- Partial differential equations (20 min)
 - Examples
 - Numerical solvers
 - PINN
- Operator learning (30 min)
 - Neural operator
 - DeepONet
 - Fourier neural operator
- Examples and extensions (30 min)
 - Weather forecast
 - And more

1. Introduction

Problems in science and engineering reduce to PDEs.









Example: heat equation

2D parabolic PDE





How long will my pan cool down?

Figure: Oleg Alexandrov

Example: heat equation



Spatial domain: $D = [0,1] \times [0,1]$ Time domain:T = [0,1]Points: $x = (x, y) = (x_1, x_2) \in D, \quad t \in T$ Solution function (Temperature) $u: T \times D \to \mathbb{R}$

Heat equation:

$$u_t = \Delta u$$

Where $\Delta u = u_{xx} + u_{yy}$ "The change in time is equal to 2nd-order difference in space."

Initial value problems

Heat equation:

$$u_t = \Delta u$$

Initial value problem: given the temperature at time t=0, What is the temperature a time t=1? Given u(0,x), what is u(1,x)?



Problems with coefficients

2D elliptic PDE:

$$-\nabla \cdot (a(x)\nabla u(x)) = f(x), \quad x \in D$$



$$\nabla u = (u_x, u_y)$$
$$\Delta u = \nabla \cdot \nabla u$$

Output: u(x)

My domain has two types of media, described by a(x)Given a(x), what is u(x)?

Numerical solver

Idea: discretize the problem onto some grid Solve a multi-variable equation of $\{u(i,j)\}, i = 0, ..., 10, j = 0, ..., 10$ Solve a linear system Au = f





Numerical solver

Idea: discretize the problem onto some grid Solve a multi-variable equation of $\{u(i,j)\}, i = 0, ..., 10, j = 0, ..., 10\}$

How to compute the derivative? Finite difference method (FDM): approximate by its neighbor points





Numerical solver

Pipeline:

- 1. Discretize the space
- 2. Write out a linear system
- 3. Solve the linear system

Common numerical methods:

- Finite difference methods
- Finite element methods
- Spectral methods
- Iterative methods

Trade-off: finer grids are more accurate, but also more expensive.

Physics-informed neural networks

Heat equation:

$$u_t = \Delta u$$

Physics-informed neural networks (PINNs)

- Idea: represent function *u* as a neural network
- Compute the derivatives $(u_t, \Delta u)$ using the chain rule
- Solve the system using gradient descent methods

PINN: <u>M Raissi</u>, <u>P Perdikaris</u>, <u>GE Karniadakis</u>

Recap: PINNs

Heat equation:

$$u_t = \Delta u$$

Pipeline:

- Initialize NN u(x)
- For N iterations
- Sample collocation points $\{(x, t)\}$
- Compute the derivatives $u_t(x,t)$, $\Delta u(x,t)$
- Minimize || $u_t \Delta u$ ||

PINN: <u>M Raissi</u>, <u>P Perdikaris</u>, <u>GE Karniadakis</u>

Recap: PINNs

Pros:

- Simple and flexible
- No need to worry about stability conditions, etc
- Complex geometry, high-dim, inverse problem, etc Cons:
- Optimization is tricky
- Less accurate than existing solvers (FDM/FEM)
- No guarantee of (optimization)

PINN: <u>M Raissi</u>, <u>P Perdikaris</u>, <u>GE Karniadakis</u>

2. Operator learning



Operator learning

Operators are map between function space. Given a dataset of input-output pairs, find the map (operator)



Operator learning

Operators are map between function space. Learn the Operators from data (coefficients & solutions pairs).

- Fix an equation
- Multiple training instances
- Learn the mapping



 $\mathcal{F}:\mathcal{A}\times\Theta\to\mathcal{U}$

Solve vs learn

Solving for a PDE instance u approximate u(x) in the spatial space.

learn the solution operator \mathcal{F} interpolate *u* in the function space.



Pipeline of operator learning

Supervised learning

Get a dataset {(a, u)} (existing solvers or experiments)

- Initialize NN, F: a->u
- For N epoch, For batches of data pairs (a, u)
- Compute the prediction F(a)
- Minimize ||F(a)-u||

Solve vs learn

Conventional methods: Solve the equation By approximation on a mesh Data-driven methods: Learn the trajectory From a distribution



Solve vs learn

Conventional methods:

- Solve for any parameters
- Worst case guarantees
- Consistency

Data-driven methods:

- Parameters from a distribution
- Less guaranteed
- Not "consistent"







Output: solutions

Architecture design

- Not vector-to-vector mapping.
- But function-to-function mapping.



Discretized vector

Continuous function

CNN vs Neural operators

Key idea: represent function & operator in mesh-invariant way



3.1 Neural operator $u = (K_l \circ \sigma_l \circ \cdots \circ \sigma_1 \circ K_0) v$

Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, Anima Anandkumar

Neural networks



$$\begin{aligned} \mathbf{v}_0 &= \mathbf{x} \\ \mathbf{v}_{l+1} &= \sigma \big(A_l \mathbf{v}_l + \mathbf{b}_l \big), \quad l = 0, \dots, L-1 \\ \mathbf{y} &= A_L \mathbf{v}_L + \mathbf{b}_L \\ \sigma : \mathbb{R} \to \mathbb{R}, \quad A_l \in \mathbb{R}^{d_{l+1} \times d_l}, \quad \mathbf{b}_l \in \mathbb{R}^{d_{l+1}} \end{aligned}$$

Neural operators



$$\begin{split} v_0(s) &= P(x(s), s) \\ v_{l+1}(s) &= \sigma \left(W_l v_l(s) + \int_D \kappa_l(s, z) v_l(z) \, dz + b_l(s) \right), \quad l = 0, \dots, L-1 \\ y(s) &= Q(v_L(s)) \\ P : \mathbb{R}^{m+n} \to \mathbb{R}^{d_0}, \quad W_l \in \mathbb{R}^{d_{l+1} \times d_l}, \quad \kappa_l : \mathbb{R}^{2n} \to \mathbb{R}^{d_{l+1} \times d_l}, \quad b_l : \mathbb{R}^n \to \mathbb{R}^{d_{l+1}}, \quad Q : \mathbb{R}^L \to \mathbb{R}^r \end{split}$$

Approximation bound

For any continuous operator, there exists a two-layers neural operators can approximate it.

For the solution operator of Navier-Stokes equation, the number of parameters depends sub-linearly on the error

Neural Operator: Learning Maps Between Function Spaces. Kovachki et. al. On universal approximation and error bounds for Fourier Neural Operators. Kovachki et. al.

Graph-based neural operators





GKN

MGKN

https://arxiv.org/abs/2006.09535 (Neurips2020)

Kernel convolution as message passing on graph

$$v_{t+1}(x) = \sigma \left(Wv_t(x) + \int_D \kappa_\phi(x, y, a(x), a(y)) v_t(y) \nu_x(dy) \right)$$

$$v_{t+1}(x) = Wv_t(x) + \sum_{y \in N(x)} \kappa_{\phi}(e(x,y))v_t(y)$$
 Graph neural network

- Adjacency matrix = kernel matrix.
- Kernel integration = message passing





3.2 Deep Operator Network



DeepONet: Lu Lu, Pengzhan Jin, George Em Karniadakis

DeepONet

F:
$$[a(x), y] \rightarrow u(y)$$

 $[(\mathbb{R}^n \rightarrow \mathbb{R}^m), \mathbb{R}^n] \rightarrow \mathbb{R}^m$

Input functions + query point \rightarrow the solution at the query point

$$(y) \longrightarrow u(y)$$
Input: a(x)

DeepONet: Lu Lu, Pengzhan Jin, George Em Karniadakis

DeepONet

Two networks:

- Branch net takes the input function
- Trunk net takes the query point



DeepONet and Neural operator

Two slightly different formulations. Potentially one can be converted into the other.

Implementation difference:

- Neural operator is easier if input and output are fully observed functions.
- DeepONet is easier when given sparse observation from sensors a(x) and query u(y).

DeepONet: Lu Lu, Pengzhan Jin, George Em Karniadakis

3.3 Fourier neural operator



Fourier neural operators



Architecture

$$v_{l+1}(s) = \sigma \left(W_l v_l(s) + \int_D \kappa_l(s,z) v_l(z) \, dz + b_l(s) \right)$$

Fourier space

- Assume $\kappa_l(s, z) = \kappa_l(s z)$ and parametrize its Fourier components θ_l .
 - Convolution theorem: $\int_D \kappa_l(s-z)v_l(z) dz = \mathcal{F}^{-1}(\theta_l \cdot \mathcal{F}(v_l))(s) \quad \mathcal{O}(n \log n)$

Fourier layer

Use convolution as the integral operator and implement with Fourier transform

$$ig(\mathcal{K}(a;\phi)v_tig)(x):=\int_D\kappaig(x,y,a(x),a(y);\phiig)v_t(y)\mathrm{d}y,$$

$$(\mathcal{K}(\phi)v_t)(x) = \mathcal{F}^{-1}(R_{\phi} \cdot (\mathcal{F}v_t))(x)$$
Fourier layer

- 1. Fourier transform
- 2. Linear transform
- 3. Inverse Fourier transform

$$(\mathcal{K}(\phi)v_t)(x) = \mathcal{F}^{-1}\Big(R_{\phi}\cdot(\mathcal{F}v_t)\Big)(x)$$



Fourier layer

The linear transform *W* outside keep the track of the location information (x) and non-periodic boundary



$$v_{t+1}(x) = \sigma \left(W v_t(x) + \int_D \kappa_\phi(x, y, a(x), a(y)) v_t(y) \nu_x(dy) \right)$$

4. Experiments



Weather Forecast

Weather forecast: the AFNO model on the ERA5 dataset

Input the current weather variables, output the weather at 6 hours later.



Pathak et al. (2022), FourCastNet: A Global Data-driven High-resolution Weather Model using Adaptive Fourier Neural Operators, arXiv: https://arxiv.org/abs/2202.11214



FNO Prediction

Ground Truth

NOAA Forecast

Actual Track

Weather Forecast

- Unprecedented skill on precipitation forecasts
- 6X higher resolution than other DL models
- 1000-member ensemble in a fraction of a second
- 4 to 5 orders of magnitude speedup over NWP
- 25000X smaller energy footprint



Pathak et al. (2022), FourCastNet: A Global Data-driven High-resolution Weather Model using Adaptive Fourier Neural Operators, arXiv: https://arxiv.org/abs/2202.11214



U-FNO CO² Storage Prediction Stanford, Caltech, Purdue, NVIDIA

Animation of Gas Saturation values over 30 years



FNO is 10⁵x faster than the FDM solver, less blurry compared to CNN

V=1e-4, zero-shot super-resolution



Bayesian inverse problem:



We use a MCMC method, sampling initial conditions and evaluating them with the traditional solver and Fourier operator. The Fourier operator takes **0.005s** to evaluate each initial condition, while the traditional solver takes **2.2s**.

Plasticity



$$\begin{aligned} \nabla \cdot S^{\varepsilon} &= \rho^{\varepsilon} u_{tt}^{\varepsilon} \\ K^{\varepsilon} &= 0 \\ u^{\varepsilon}(x,0) &= u_0(x), \quad u_t^{\varepsilon}(x,0) = v_0(x), \quad \xi^{\varepsilon}(x,0) = \xi_0 \\ u^{\varepsilon}(x,t) &= u^*(x,t) \\ S^{\varepsilon}(\nabla u^{\varepsilon},\xi^{\varepsilon},x)n(x) &= s^*(x,t) \end{aligned}$$

Multi-scale method: use neural operator to map from strain to stress on the unit cell; update macroscale with Abaqus solver.

Plasticity



PCA-operator solves multi-scale plasticity problem (Burigede et. al.)

Ultrasound



lithography





Benchmark	UNet@10 mPA (%)	epoch [25] mIOU (%)	DAMO-DI mPA (%)	S@10epoch [10] mIOU (%)	DAMO-DL mPA (%)	S@100epoch [10] mIOU (%)	Ours@ mPA (%)	10epoch mIOU (%)
ISPD-2019 (L)	99.40	98.03	98.40	97.50	99.25	98.11	99.43	98.27
ISPD-2019 (H)	99.08	97.97	-	-	-	-	99.21	98.45
ICCAD-2013 (L)	97.30	95.38	96.24	95.15	98.94	96.97	98.98	97.79
ICCAD-2013 (H)	95.16	93.04	-	-	-	-	99.12	97.77

Haoyu Yang et. al.

5. Extensions



Zongyi Li*, Hongkai Zheng*, Nikola Kovachki, David Jin, Haoxuan Chen, Burigede Liu, Kamyar Azizzadenesheli, Anima Anandkumar

Data loss: compare prediction and ground-truth solution



Equation loss: plug prediction in PDE and compute residual

$$-\nabla \cdot (a(x)\nabla u(x)) = f(x), \quad x \in D$$
$$u(x) = 0, \qquad x \in \partial D$$

PINO: PHYSICS-INFORMED NEURAL OPERATOR

Neural operators use only the data. If the explicit form of the PDE is known, we can combine neural operators with the PINNs setting.

- 1. Pre-train
- operator-learning setting
- a family of PDE
- use the data loss (and equation loss)



- 2. test-time optimize
- solver setting (PINN)
- a specific instance
- use the equation loss



Standard cases (Burgers, Darcy, KF, Cavity): no need to use any data

• Pretrain improves both the convergence speed and final error.

•



- PINO gets 2% error on Re500 [$2\pi \times 2\pi \times 1s$]
- Easily generalize from one Re to another



Relative error: PINO: 0.9%, PINN:18.7%

- Pretrain improves both the convergence speed and final error.
- More robust to hyperparameters



If data is available: PINO can solve very challenging scenarios such as long-temporal transient flow T=50, which is intractable to PINNs.



Long-temporal transient flow: preserve the 400x speedup compared to the pseudospectral solver.



Solving inverse problems



Backward PINO accurately solves inverse problem. 3000x speedup

Transfer Learning





Pre-train on Re100, fine-tune on Re500. Converge 3x faster.

Geometric-FNO

Zongyi Li, Zhengyu Huang, Burigede Liu, Anima Anandkumar

Geometric FNO

Previous FNO implementation relies on FFT. It can be only applied to rectangular domains with a uniform grid. We want to extend FNO to different geometries.

Further, we want to learn solution operators mapping from the shapes to the solution, and solve the inverse design problems.





Geometric FNO

- Any boundary -> embed (Fourier continuation)
- Any mesh -> learned interpolation /DFT
- Any shape -> deform (adaptive/moving mesh)
- Any topology -> decompose (pants/handler decomposition)

boundary

For non-periodic boundary, we can embed the domain into a larger domain with periodic condition



Similar, for any obstacles and holes. we can fill in zeros.





Mesh



Given any input mesh, we first transform it to a uniform mesh.

We use the learned interpolation /discrete Fourier transform to reduce the interpolation error.

Shape





Given any input domain, we want to find a deformation (diffeomorphism) To convert the domain into a regular one.

Adaptive meshes



Such deformation can be also used to construct adaptive meshes for multiscale structures

Examples: elastic equation

A constant forcing is applied from the left. Given the shape, we want to predict the stress. Unstructured mesh + irregular shape.



Examples: elastic equation



Examples: elastic equation

Using adaptive meshes further improves the performance.





	Train Raw	Test Raw	Test DFT
Uniform	7%	10%	11%
X-Y deform	9%	10%	9%
Theta-R	6%	6%	~6%
Examples: Channel flow (pipe)



Examples: Airfoils



Examples: inverse design

Optimize the spline nodes to achieve the design objectives.

Pipe: optimize the upper (and lower) boundary









PDE-observer

Yuanyuan Shi, Zongyi Li, Huan Yu, Drew Steeves, Anima Anandkumar, and Miroslav Krstic

MOTIVATING PROBLEM IN ROBOTICS





Given sensors placed on the body of airfoil/drone (boundary), Can we estimate the velocity field (interior) for more accurate control?

Yuanyuan Shi, Zongyi Li, Huan Yu, Drew Steeves, Anima Anandkumar, and Miroslav Krstic

STATE ESTIMATION THROUGH PDE OBSERVER

- Given the partial observation, estimate the state.
- PDE for state estimation: Back-stepping PDE observer.
- The estimate converges to the truth (exponentially)
- Slow and not real-time.



Neural observer with FNO

Recurrent observer: make the estimation at each step recursively



Feedforward observer: make the estimate over a time interval each step.



REACTION-DIFFUSION PDE PRESCRIBED-TIME OBSERVER



Chemical tubular reactor system: Learn a time-dependent, fast-converging observer.





2000x Speed up

FNO-transformer

John Guibas , Morteza Mardani , Zongyi Li , Andrew Tao, Anima Aanandkumar, Bryan Catanzaro

Fourier-based Transformer

Neural operator can be viewed as a continuous generalization of Transformers

- Attention mechanism is a kernel-integration.
- Query-key can be viewed a a low-rank decomposition of the kernel.
- Replace the kernel by Fourier transform.



Shuhao Cao. Choose a Transformer: Fourier or Galerkin

Fourier-based Transformers



Backbone	Mixer	Params	GFLOPs	Latency(sec)	SSIM	PSNR(dB)
ViT-B/4	Self-Attention	87M	357.2	1.2	0.931	27.06
ViT-B/4	LS	87M	274.2	1.4	0.920	26.18
ViT-B/4	GFN	87M	177.8	0.7	0.928	26.76
ViT-B/4	AFNO (ours)	87M	257.2	0.8	0.931	27.05



Chaotic system

Miguel Liu-Schiaffini, Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, Anima Anandkumar

Chaotic system

Chaotic systems are intrinsically unstable. Smaller errors will accumulate and make the simulation diverge from the truth.

Can we predict long-time trajectories that, while eventually diverging from the truth, still preserve the same orbit (attractor) of the system and its statistical properties?



Semigroup and Markov neural operator

Use Markov neural operator to model the local evolution.

$$u(t) = S_t u(0).$$

Compose the operator as a semigroup

$$u(n \cdot h) \approx \hat{S}_h^n(u_0) \coloneqq \underbrace{(\hat{S}_h \circ \dots \circ \hat{S}_h)}_{n \text{ times}}(u_0)$$

KS equation



FNO captures the trajectory of chaotic systems for a longer period compared to RNNs.

Markov neural operator captures the invariant statistics



Markov neural operator simulates the global attractor



Enforcing dissipativity: Motivations

Many chaotic systems have the **dissipative** property, which pushes the dynamics back to the attractor.

Without such dissipativity, the model is prone to blow up.



Accelerating flow

Methods for Enforcing dissipativity

Enforce dissipativity by augmenting virtual data points around the raw data (equivalent to adding a loss term)



$${\cal L}_{
m data} + rac{lpha}{2}\int |\Phi(x, heta)-\lambda x|^2 \
u(dx)$$

Red: the raw data (attractor) Blue shell: the region to add data points.

Lorenz 63 learned flow maps



Red points are training points on attractor, blue circles represents shell where dissipativity is enforced.

Kolmogorov flow with enforced dissipativity

Start with an initial condition outside the attractor



Before: quickly blowing-up (10¹⁶)



After: converge back to the attractor

Kolmogorov flow with enforced dissipativity



The distribution of dissipation is improved

Self-similar flow

Haydn, Zongyi Li, Yixuan Wang, Thomas Hou, Anima Anandkumar

The existence of the solution to the Navier-Stokes equation is one of the most interesting problems in PDE.

Can we construct a counterexample (blowup)?



Idea: find a blowup scenario with a self-similar structure, such as **u**->inf as t->T.

$$u(t,x) = c_0(T-t)^{-c_1}U(x \cdot (T-t)^{c_2}).$$

The original solution \mathbf{u} is singular, but the profile function \mathbf{U} is smooth. Need to find the profile parameters c1 and c2 (inverse problem).

We can use the physics-informed methods (PINN and PINO) to discover such profiles.

Pros: can find the profile parameters (inverse problem) Cons: less accurate (10⁻⁵) compared to the conventional solver (10⁻¹⁰, adaptive FEM). Since the Euler equation is highly sensitive, extreme accuracy is needed.



Y. Wang, C.-Y. Lai, J. G'omez-Serrano, and T. Buckmaster. Self-similar blow-up profile for the boussinesq equations via a physics-informed neural network.

Use PINO to improve the optimization landscape:

- High frequency (singular) structures are challenging for neural networks
- Use the adaptive PINO (like the adaptive FEM) to better capture the singularity



Reference

Arxiv:

https://arxiv.org/abs/2003.03485 https://arxiv.org/abs/2006.09535 https://arxiv.org/abs/2010.08895

Code:

https://github.com/zongyi-li/graph-pde https://github.com/zongyi-li/fourier_neural_operator

Blog posts:

https://zongyi-li.github.io/blog/2020/graph-pde/ https://zongyi-li.github.io/blog/2020/fourier-pde/