

A Roadmap to Practical Neural PDE Solvers

Mingsheng Long

School of Software, Tsinghua University

July 2024



清華大學
Tsinghua University



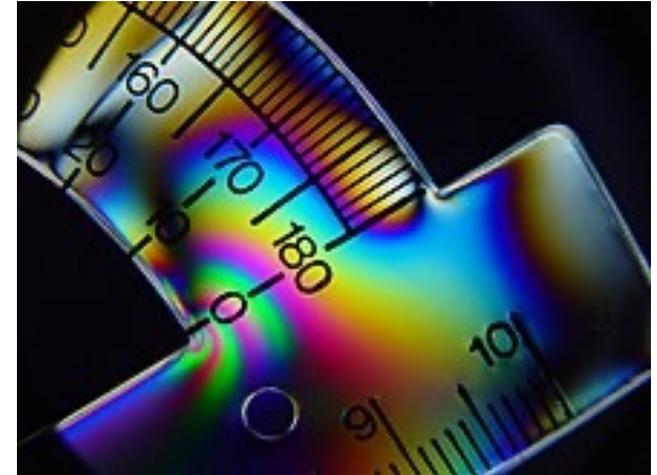
Real-world phenomena



Turbulence



Atmospheric circulation



Stress

How to understand the world?

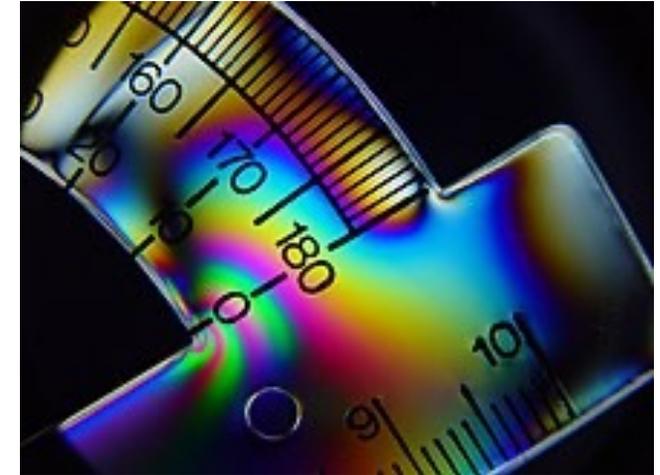
Real-world phenomena



Turbulence



Atmospheric circulation



Stress

How to understand the world?

Images? Videos?

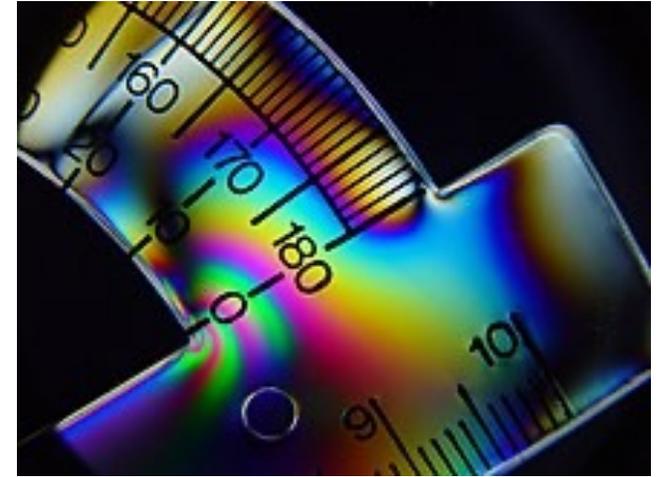
Real-world phenomena



Turbulence



Atmospheric circulation



Stress

Beyond appearances, these phenomena are governed by **scientific rules.**

Partial Differential Equations (PDEs)

➤ Fluid physics:

Navier-Stokes Equation
for fluid dynamics

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{U}) = 0$$

$$\frac{\partial \mathbf{U}}{\partial t} + \mathbf{U} \cdot \nabla \mathbf{U} = \mathbf{f} + \frac{1}{\rho} \nabla \cdot (\mathbf{T}_{ij} \mathbf{e}_i \mathbf{e}_j)$$

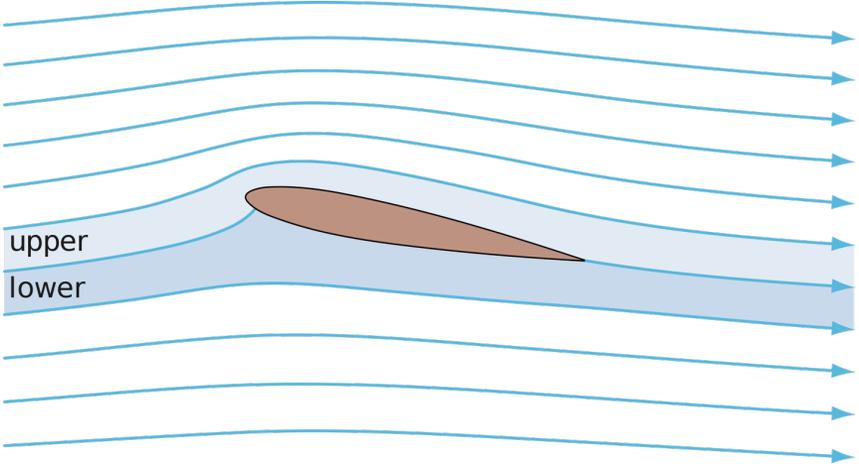
$$\frac{\partial (e + \frac{1}{2} \mathbf{U}^2)}{\partial t} + \mathbf{U} \cdot \nabla (e + \frac{1}{2} \mathbf{U}^2) = \mathbf{f} \cdot \mathbf{U} + \frac{1}{\rho} \nabla \cdot (\mathbf{U} \cdot \mathbf{T}_{ij} \mathbf{e}_i \mathbf{e}_j) + \frac{\lambda}{\rho} \Delta T$$

➤ Solid physics:

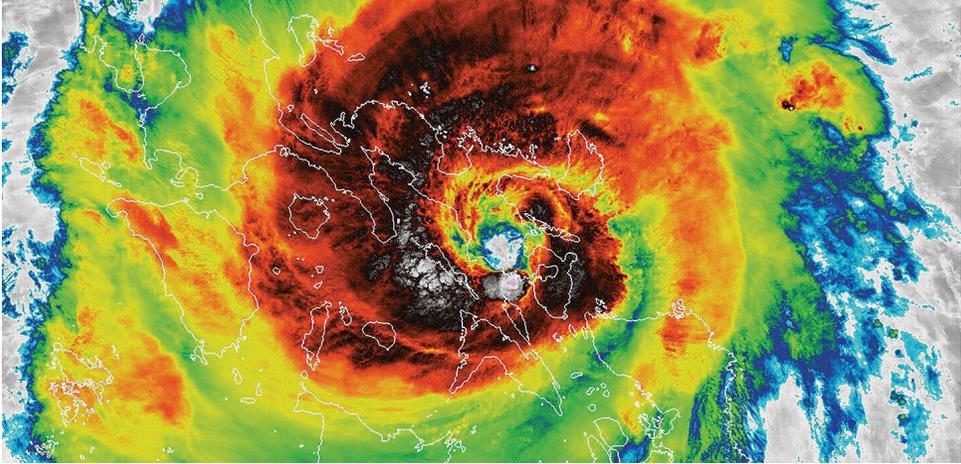
$$\rho^s \frac{\partial^2 \mathbf{u}}{\partial t^2} + \nabla \cdot \boldsymbol{\sigma} = 0$$

Inner stress
of solid materials

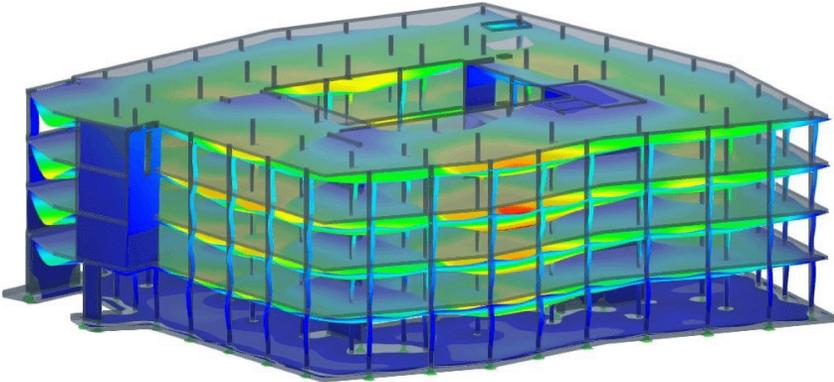
Wide Applications



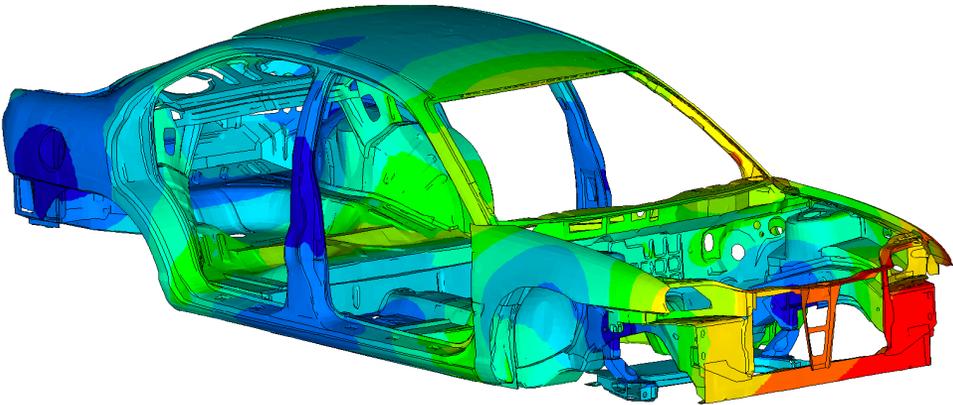
Airfoil design



Weather forecasting



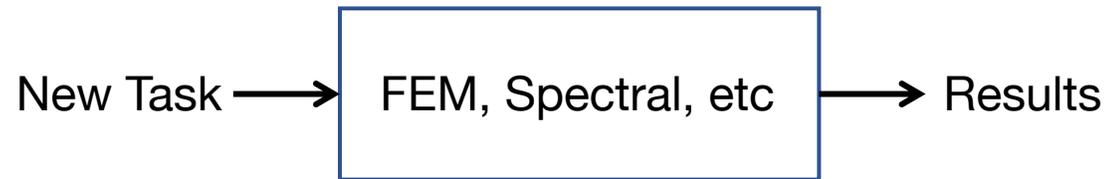
Civil engineering



Vehicle manufacturing

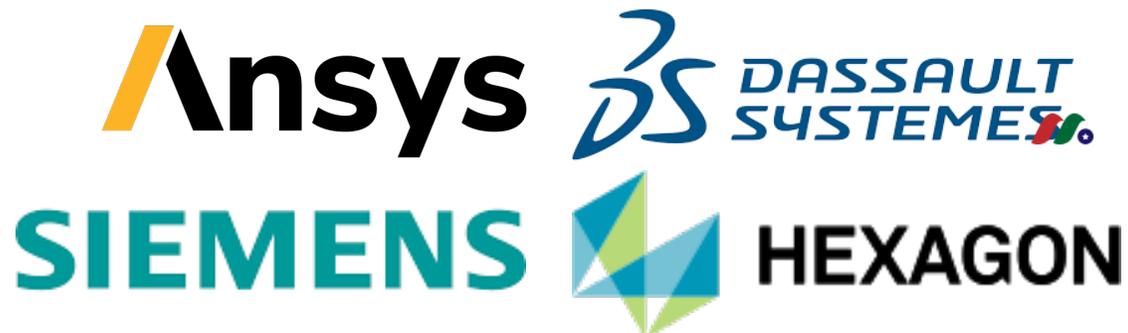
PDE Solvers

Classic Numerical Methods



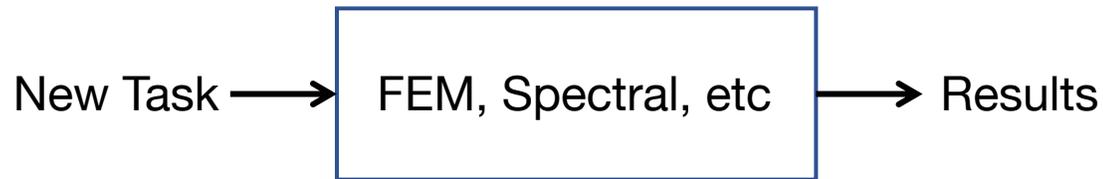
- Recalculation for every new sample
- Each round will take hours or even days

Stable but Slow



PDE Solvers

Classic Numerical Methods

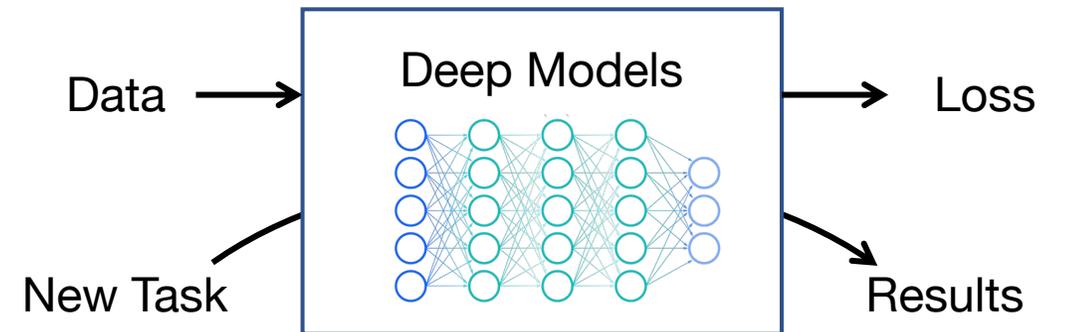


- Recalculation for every new sample
- Each round will take hours or even days

Stable but Slow



Neural PDE Solver

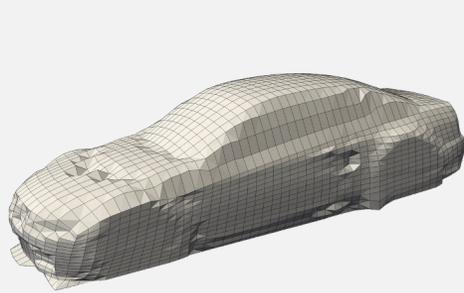


- Training once, inference a lot
- Each round needs several seconds

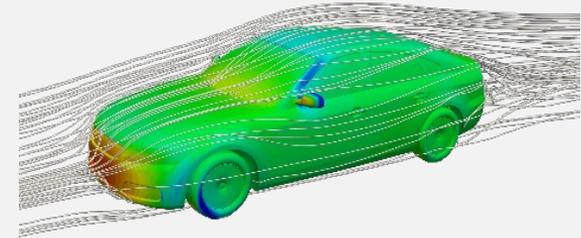
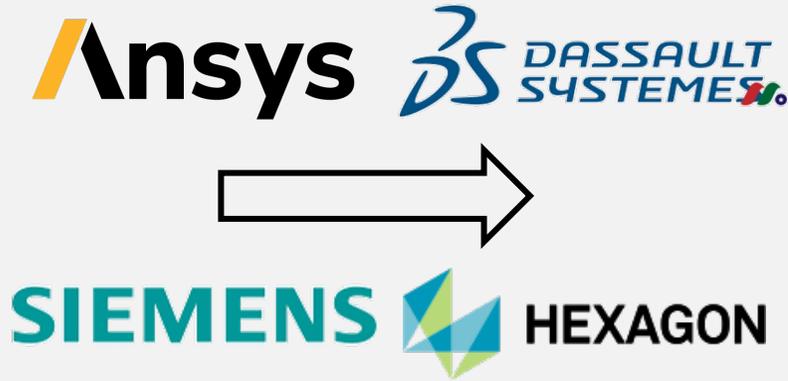
**An efficient surrogate tool
(In expectation)**

A Roadmap to Practical Neural PDE Solvers

Industrial simulation with CAE



Varied Geometries



Physical Simulation

Neural PDE Solver (Our work)

Q1: High-dimensional

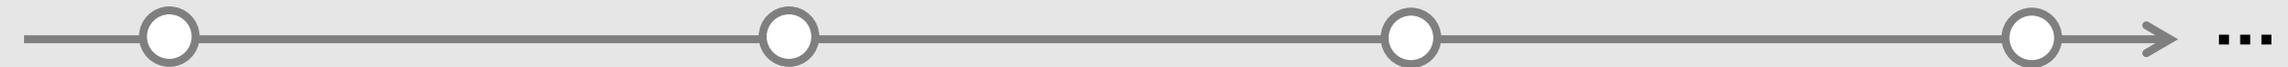
Mapping Approximation

Q2: Large-scale

Irregular Meshes

Q3: Generalization

among varied PDEs



Deep Models

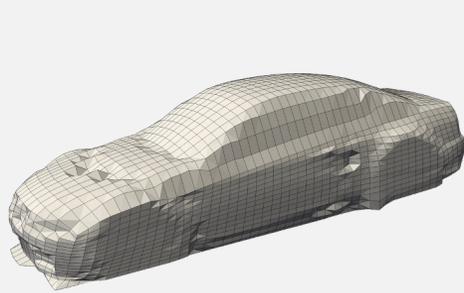
Latent Spectral Models

Transolver

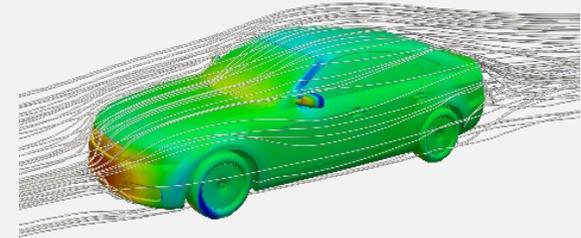
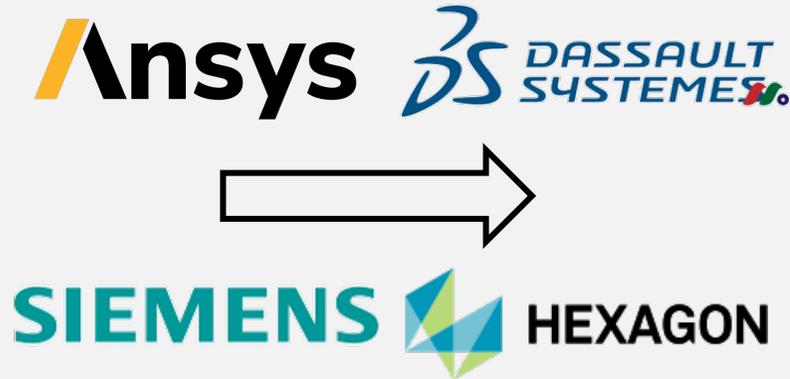
Unisolver

A Roadmap to Practical Neural PDE Solvers

Industrial simulation with CAE



Varied Geometries



Physical Simulation

Neural PDE Solver (Our work)

Q1: High-dimensional

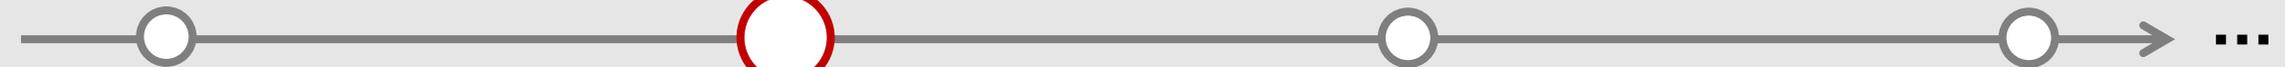
Mapping Approximation

Q2: Large-scale

Irregular Meshes

Q3: Generalization

among varied PDEs



Deep Models

Latent Spectral Models

Transolver

Unisolver



ICML | 2023

The Fortieth International Conference on Machine Learning



Solving High-Dimensional PDEs with Latent Spectral Models

Haixu Wu¹ Tengge Hu¹ Huakun Luo¹ Jianmin Wang¹ Mingsheng Long¹



Haixu Wu



Tengge Hu



Huakun Luo



Jianmin Wang



Mingsheng Long

Solving PDEs: Discretization

Infinite-dimensional
PDE solutions

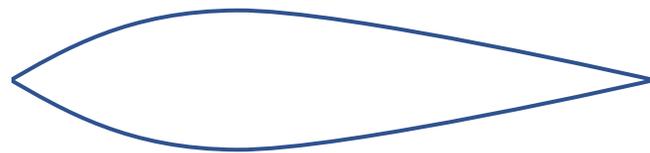
$$\mathbf{x}(\mathbf{s}), \mathbf{s} \in \mathcal{D}$$

Discretization

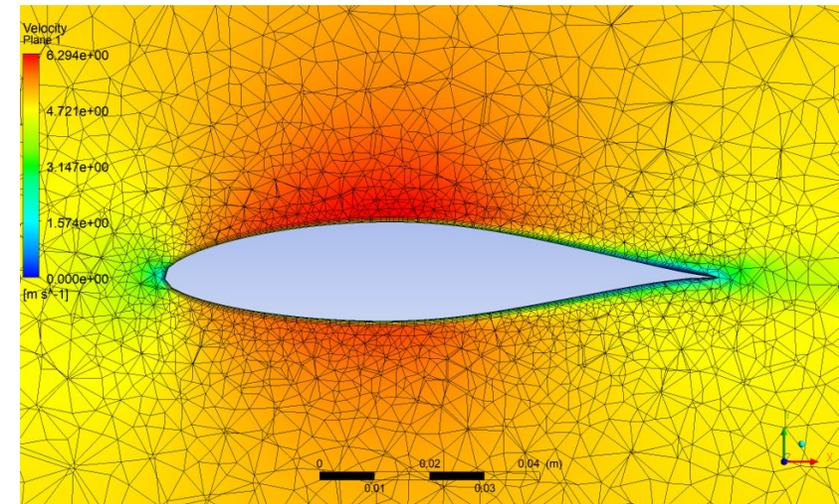
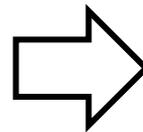


High-dimensional
coordinate spaces

\mathcal{D} is the mesh point set



Spatial continuous



Solving PDEs: Discretization

Infinite-dimensional
PDE solutions

$$\mathbf{x}(\mathbf{s}), \mathbf{s} \in \mathcal{D}$$

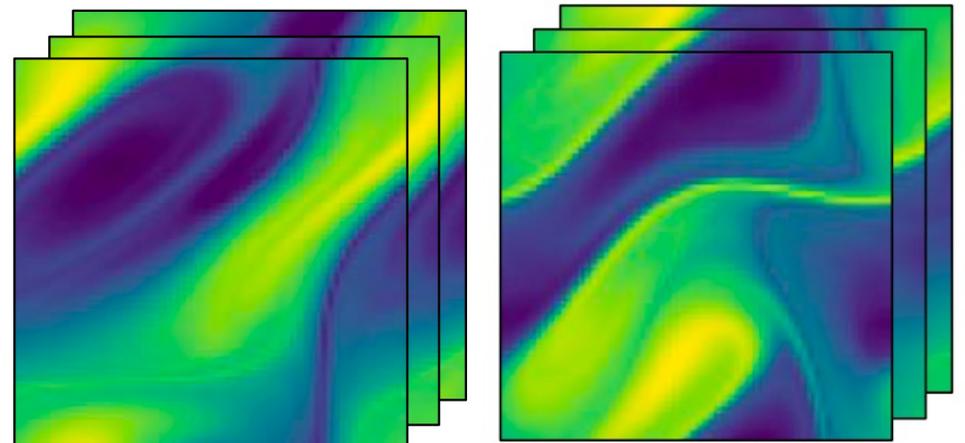
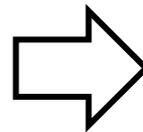
Discretization



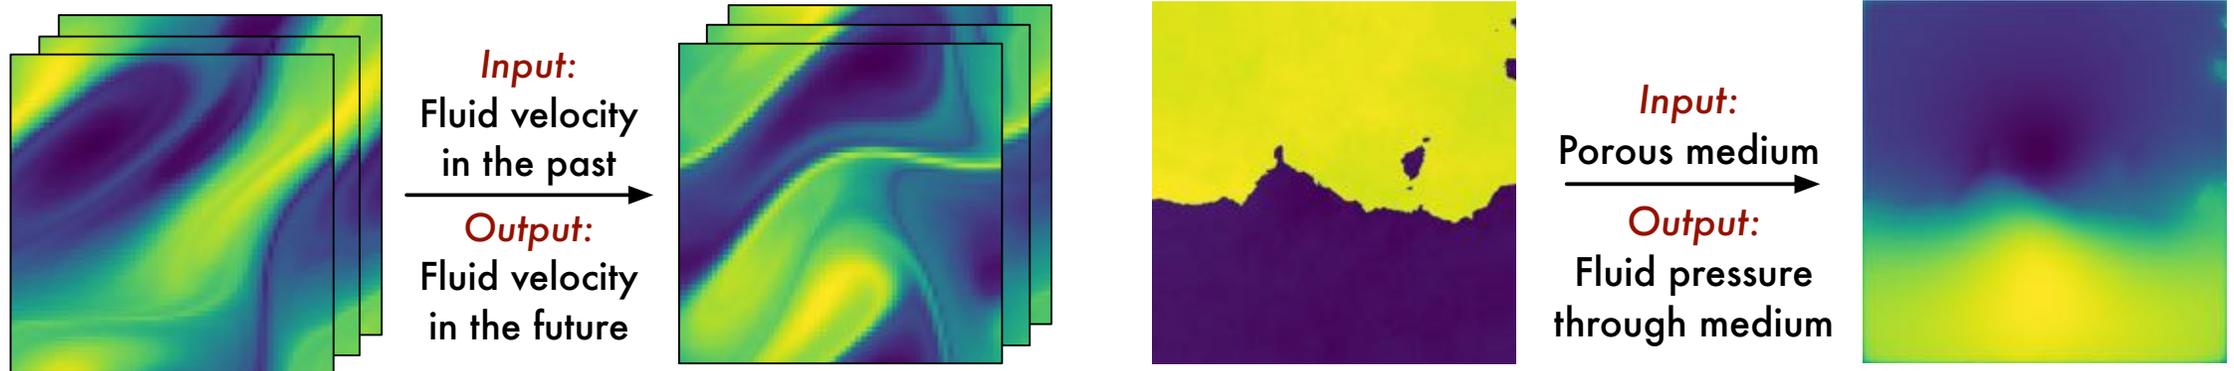
High-dimensional
coordinate spaces

\mathcal{D} is the grid point set

Spatiotemporal Continuous
Navier-Stokes Equation

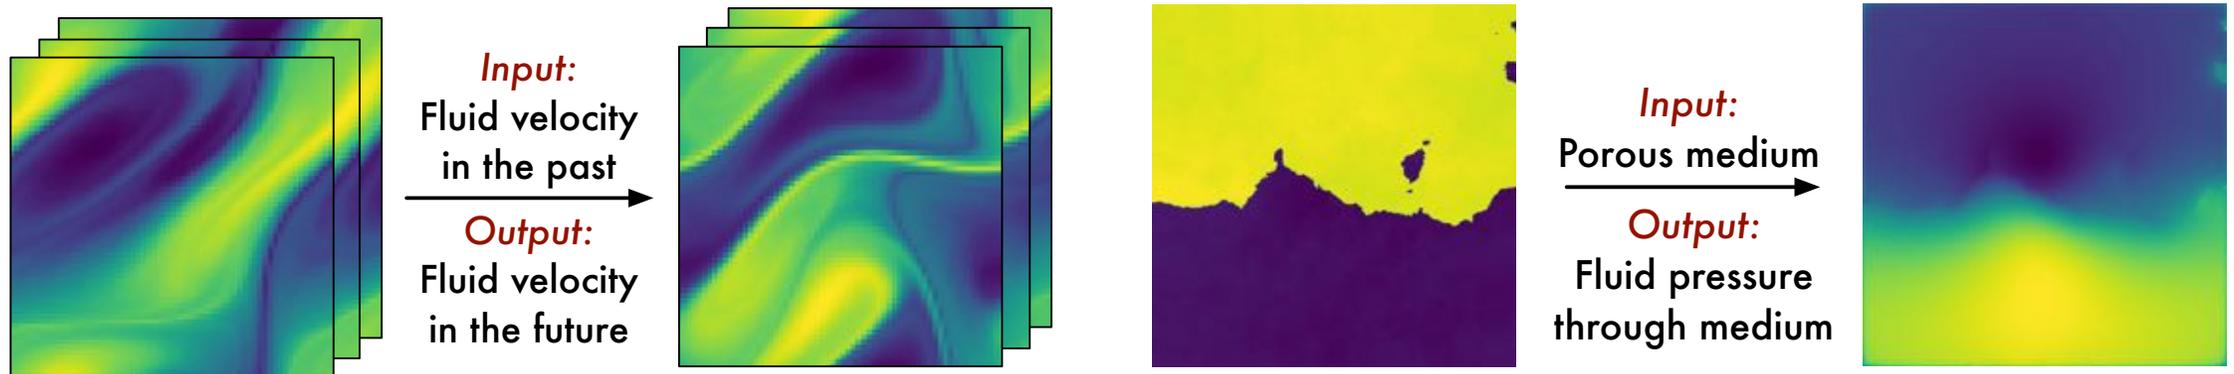


Challenges in Solving High-dimensional PDEs



- Curse of dimensionality → Huge computation cost
- Intricate interactions among physical variates of coupled equations → Complex mappings

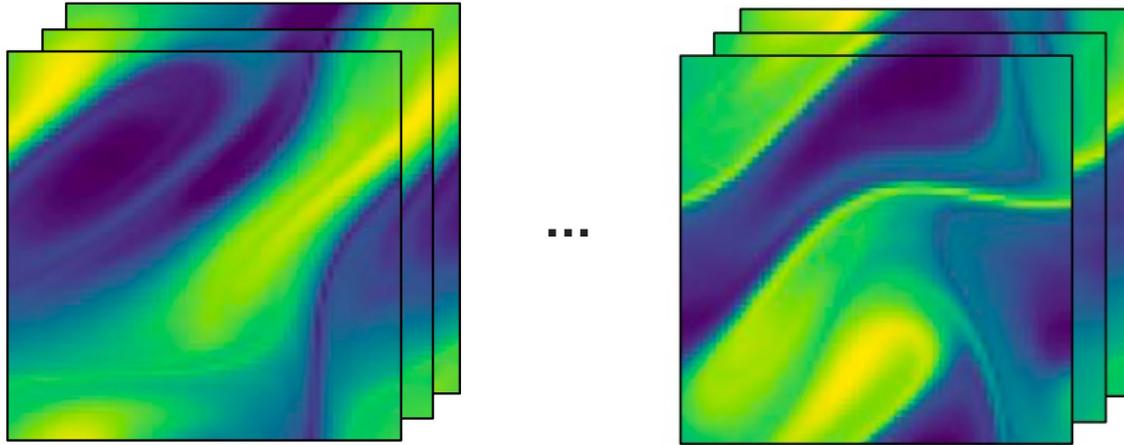
Challenges in Solving High-dimensional PDEs



- Curse of dimensionality → Huge computation cost
- Intricate interactions among physical variates of coupled equations → Complex mappings

*How to efficiently and precisely approximate complex mappings
between high-dimensional input-output pairs?*

Motivation

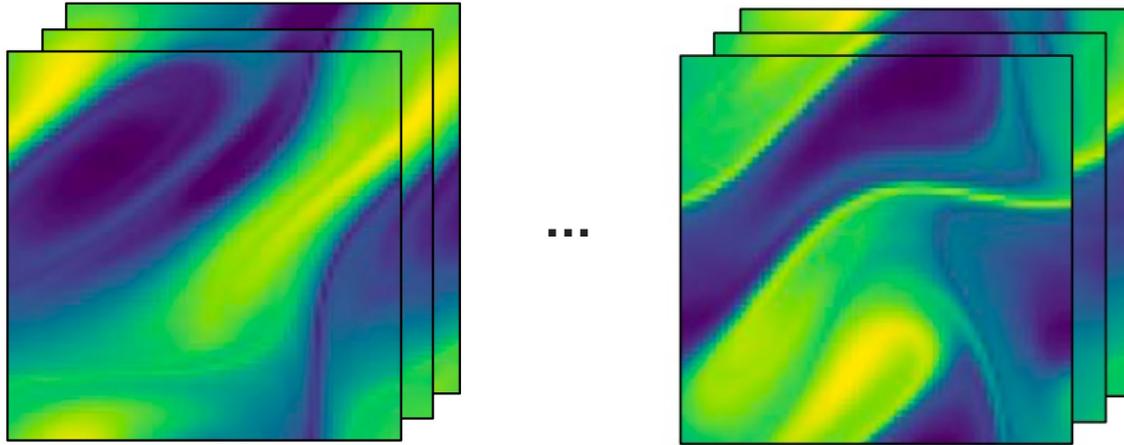


Multitudinous Data

Following the same PDE constraint

Manifold Hypothesis: *Real-world high-dimensional data lie on low-dimensional manifolds embedded within the high-dimensional space.*

Motivation



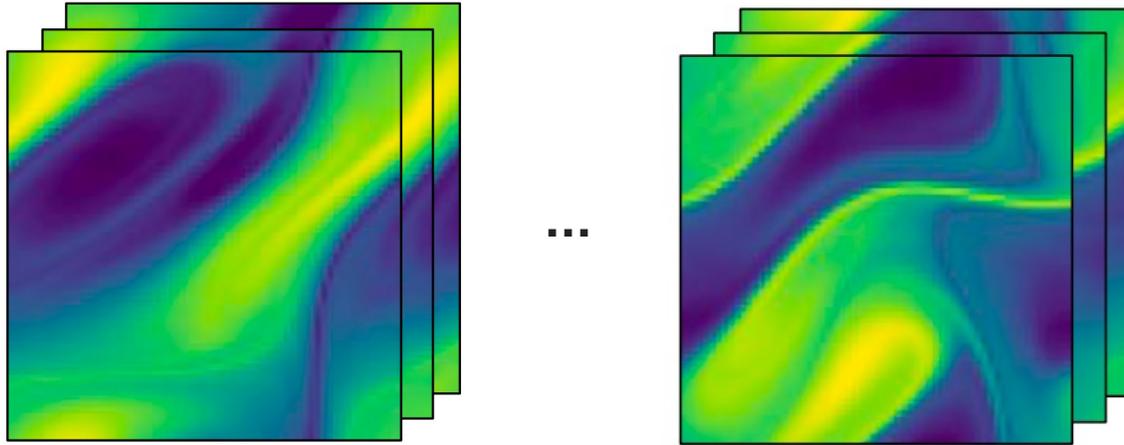
Multitudinous Data

Following the same PDE constraint

Manifold Hypothesis: *Real-world high-dimensional data lie on low-dimensional manifolds embedded within the high-dimensional space.*

1. High-dimensional data can be projected to a more compact latent space

Motivation



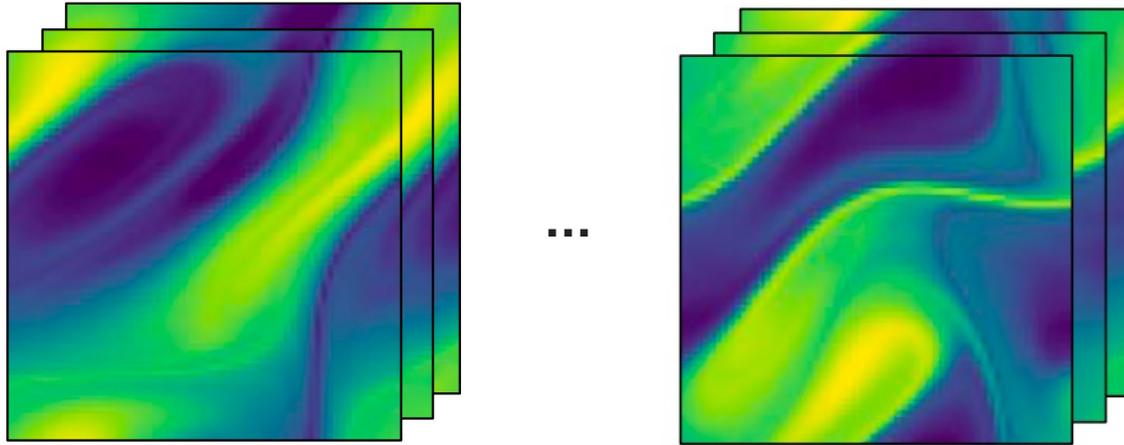
How to approximate
complex input-output mappings?



Previous Methods: Directly approximating with a single deep model

Suffer from optimization problem and limited performance

Motivation

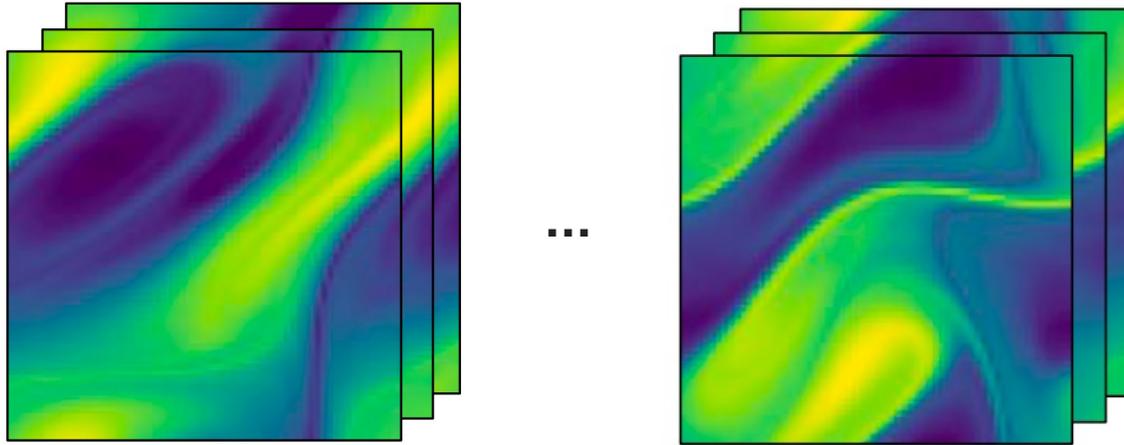


How to approximate
complex input-output mappings?



Spectral Methods: approximate solution f of a certain PDE as a *finite sum* of N orthogonal basis functions $\{f_1, f_2, \dots, f_N\}$, that is: $f \approx f^N = \sum_{i=1}^N w_i f_i$.

Motivation



How to approximate
complex input-output mappings?



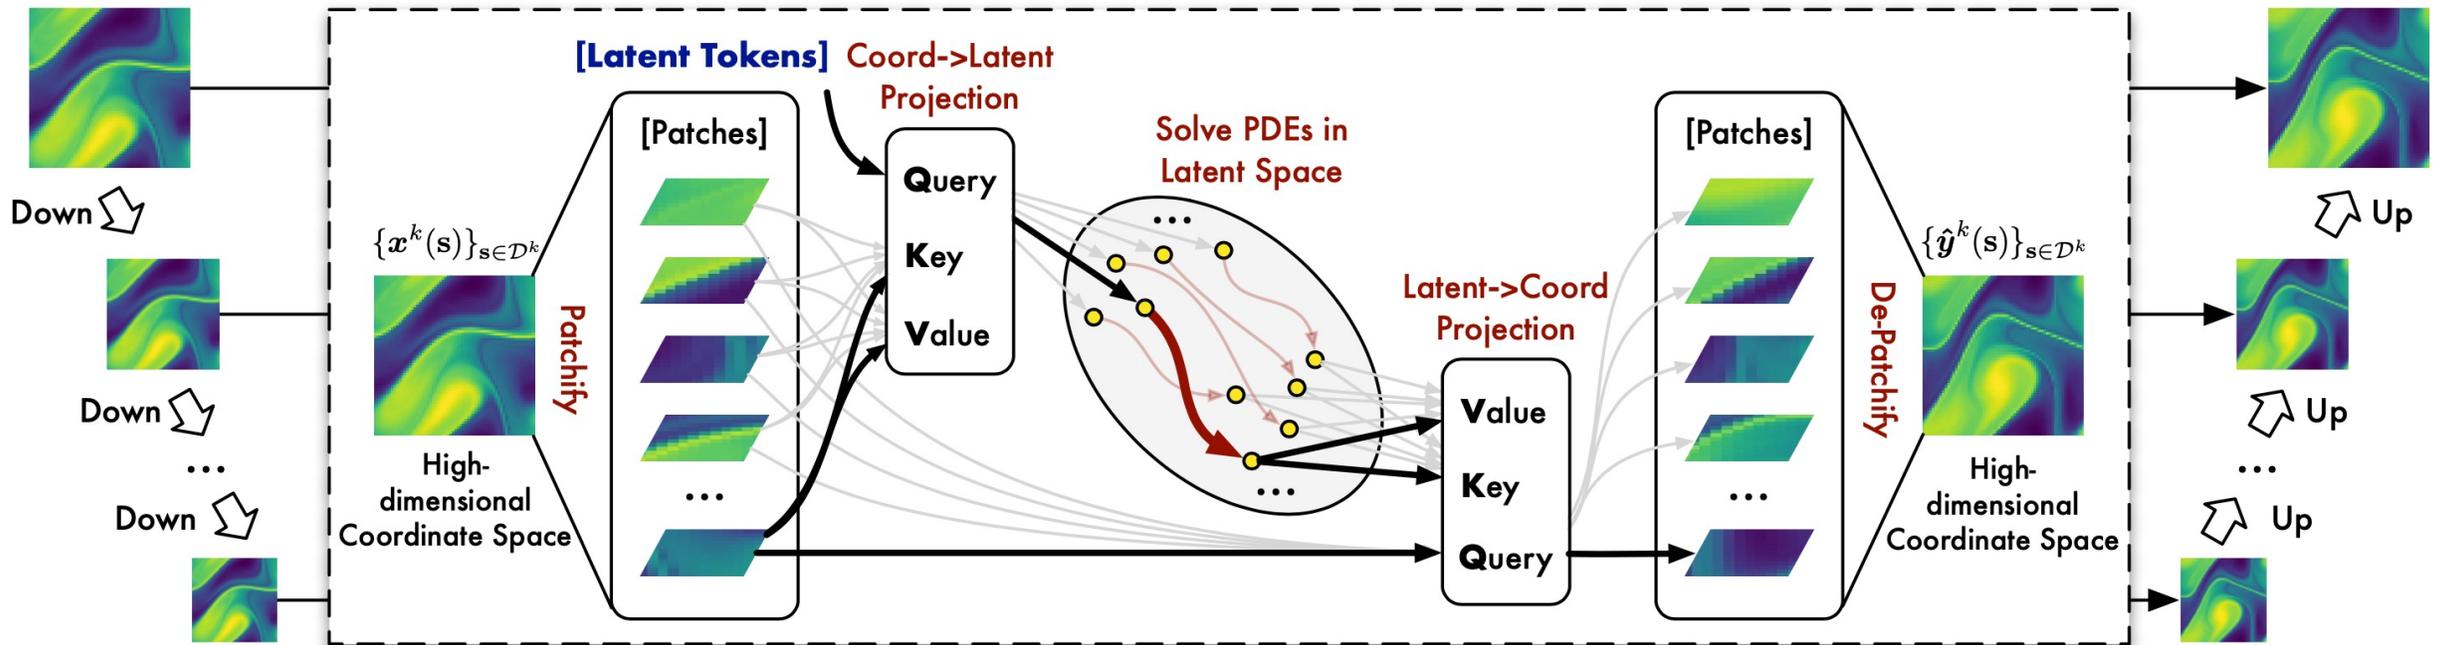
Spectral Methods: approximate solution f of a certain PDE as a *finite sum* of N orthogonal basis functions $\{f_1, f_2, \dots, f_N\}$, that is: $f \approx f^N = \sum_{i=1}^N w_i f_i$.

2. Learning multiple basis operators for approximation

Latent Spectral Models (LSM)

	Previous Methods	LSM (ours)
Solving Process	<p>Solving in the coordinate space</p> <ul style="list-style-type: none">• Huge computation cost• Making input-output mappings extremely complex	<p>Solving in the latent space</p> <ul style="list-style-type: none">• Efficient computation• Highlight the inherent physics properties
Mapping approximation	<p>Directly learning a single operator</p> <ul style="list-style-type: none">• Fail in approximating complex mappings• Lack of theoretical guarantee	<p>Learning multiple basis operators</p> <ul style="list-style-type: none">• Nice approximating and convergence properties under theoretical guarantee

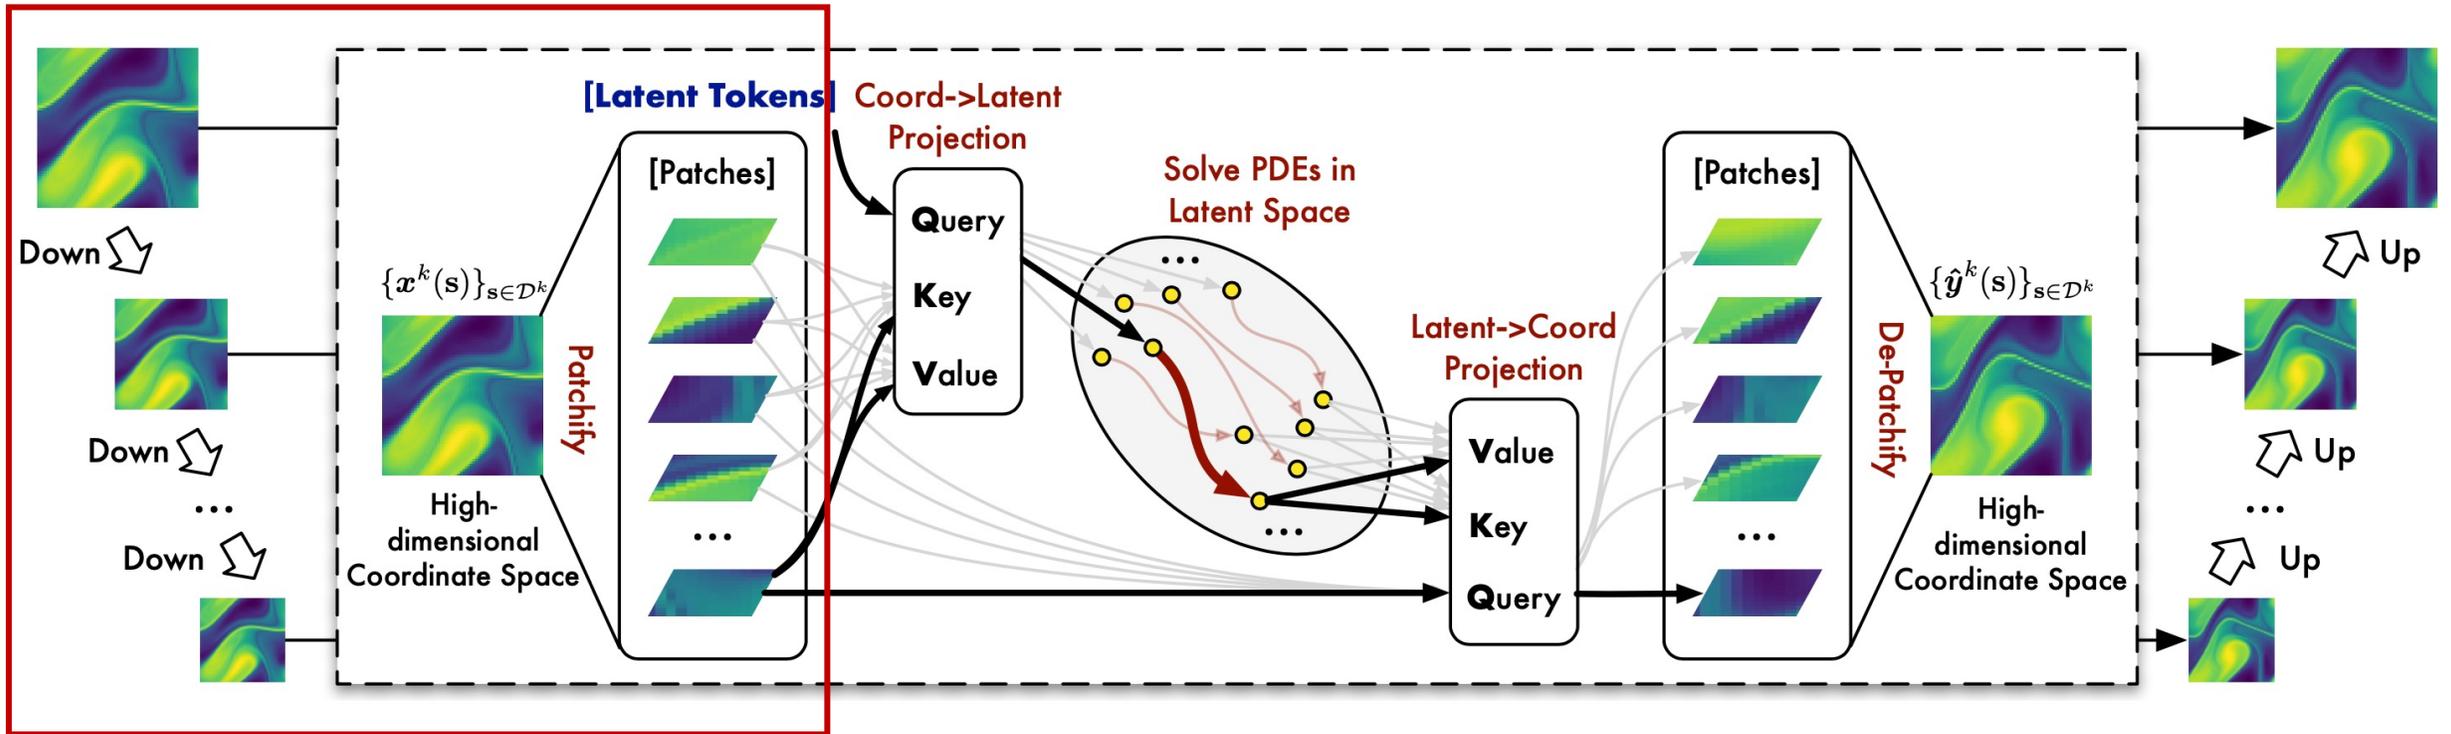
Overall design of LSM



LSM with *Hierarchical Projection Network* and *Neural Spectral Block*

① Coor \rightarrow Latent ② Solving in the Latent Space ③ Latent \rightarrow Coor

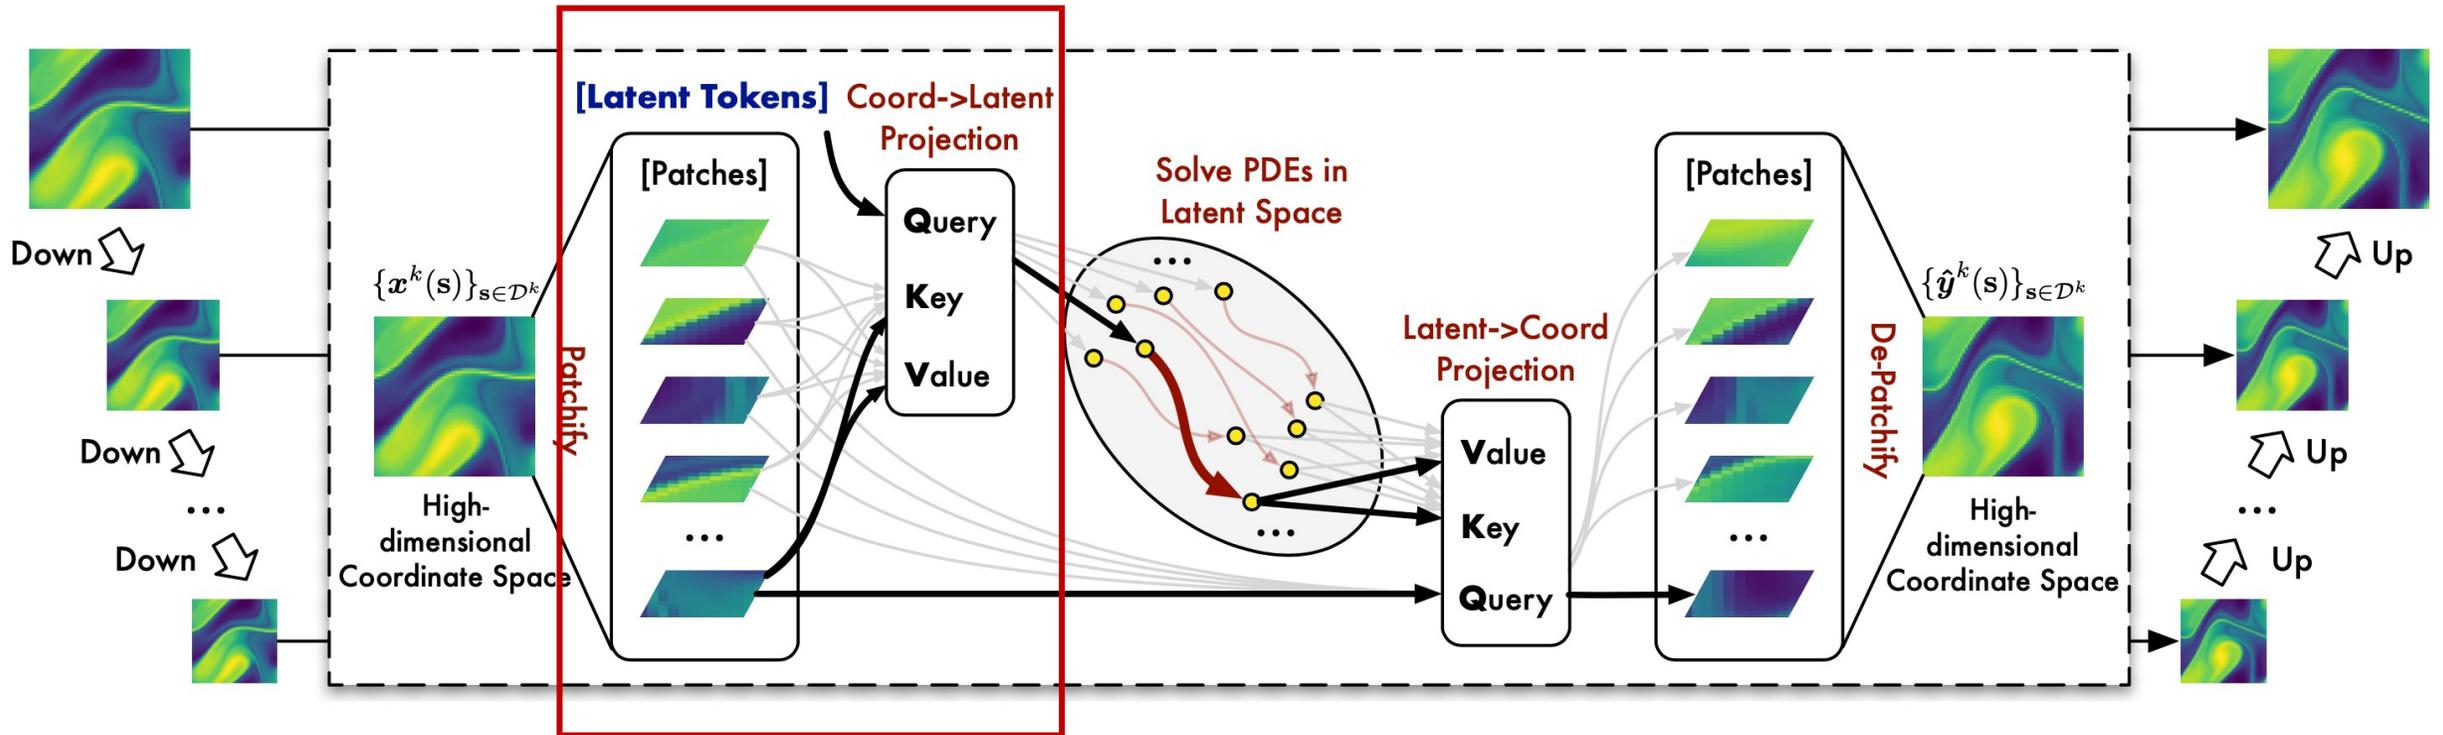
Hierarchical Projection Network



① **Multiscale patchified architecture** → Solve PDEs in different regions and scales

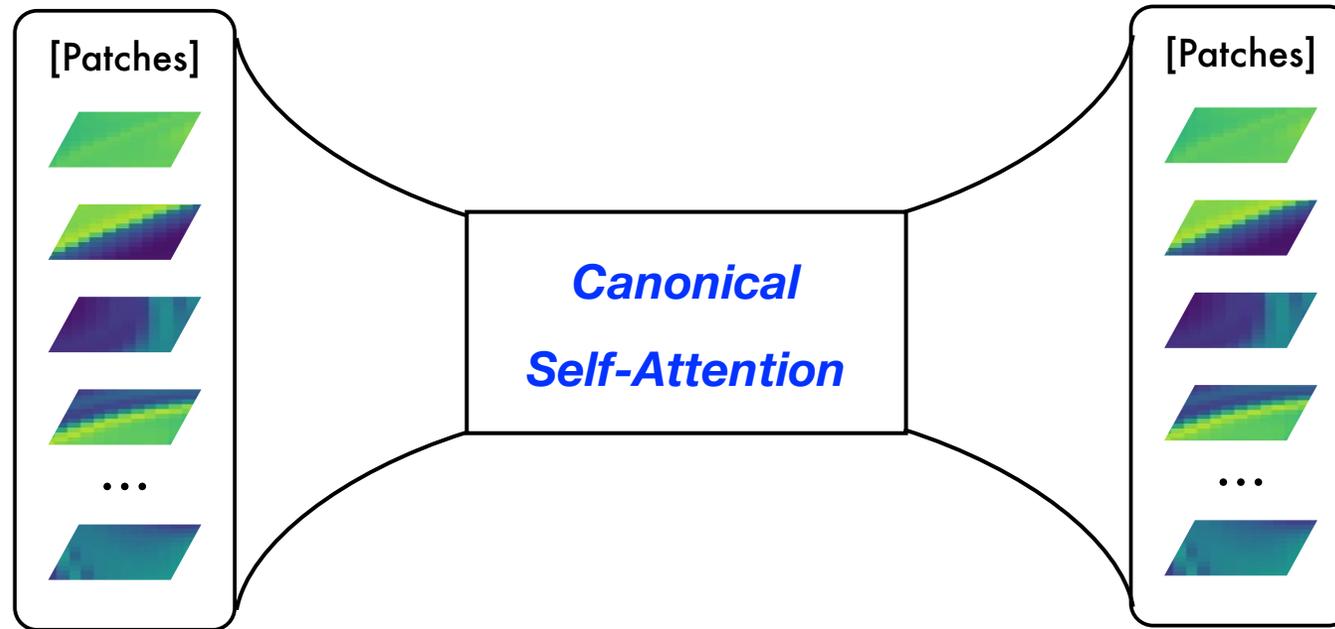
PDEs always present different physical states according to the observed scales and regions.

Hierarchical Projection Network



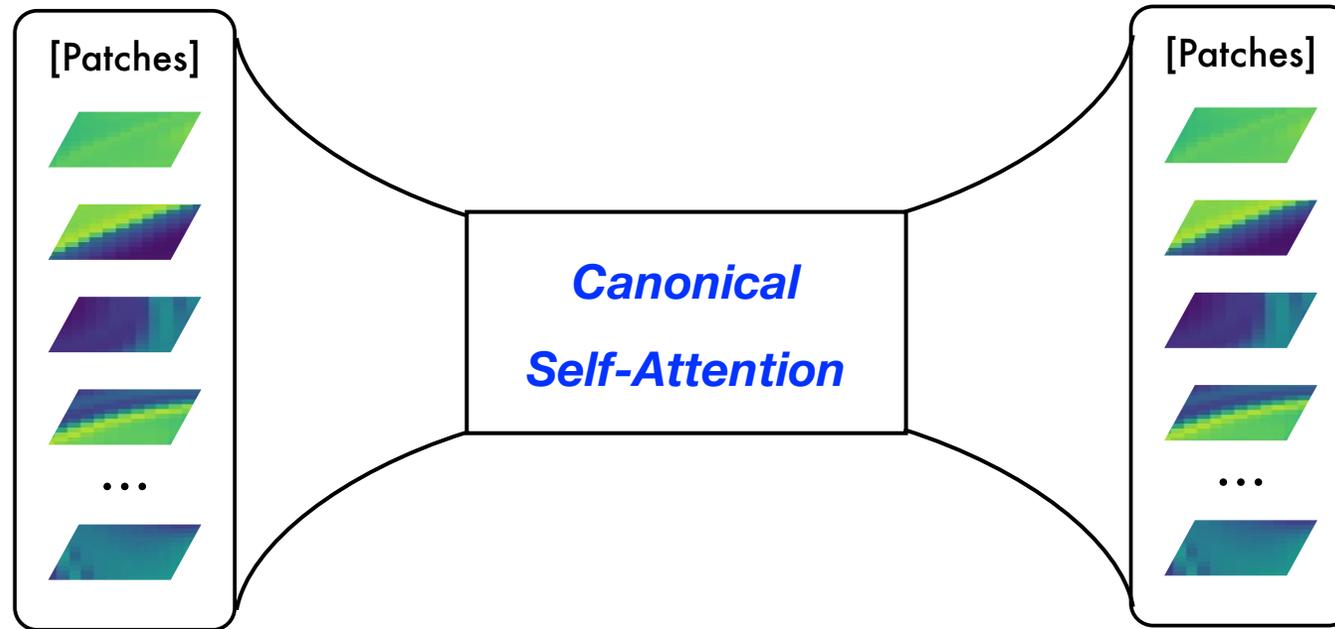
② **Attention-based projector** → Remove unwieldy coordinate information

Hierarchical Projection Network



② **Attention-based projector** → Remove unwieldy coordinate information

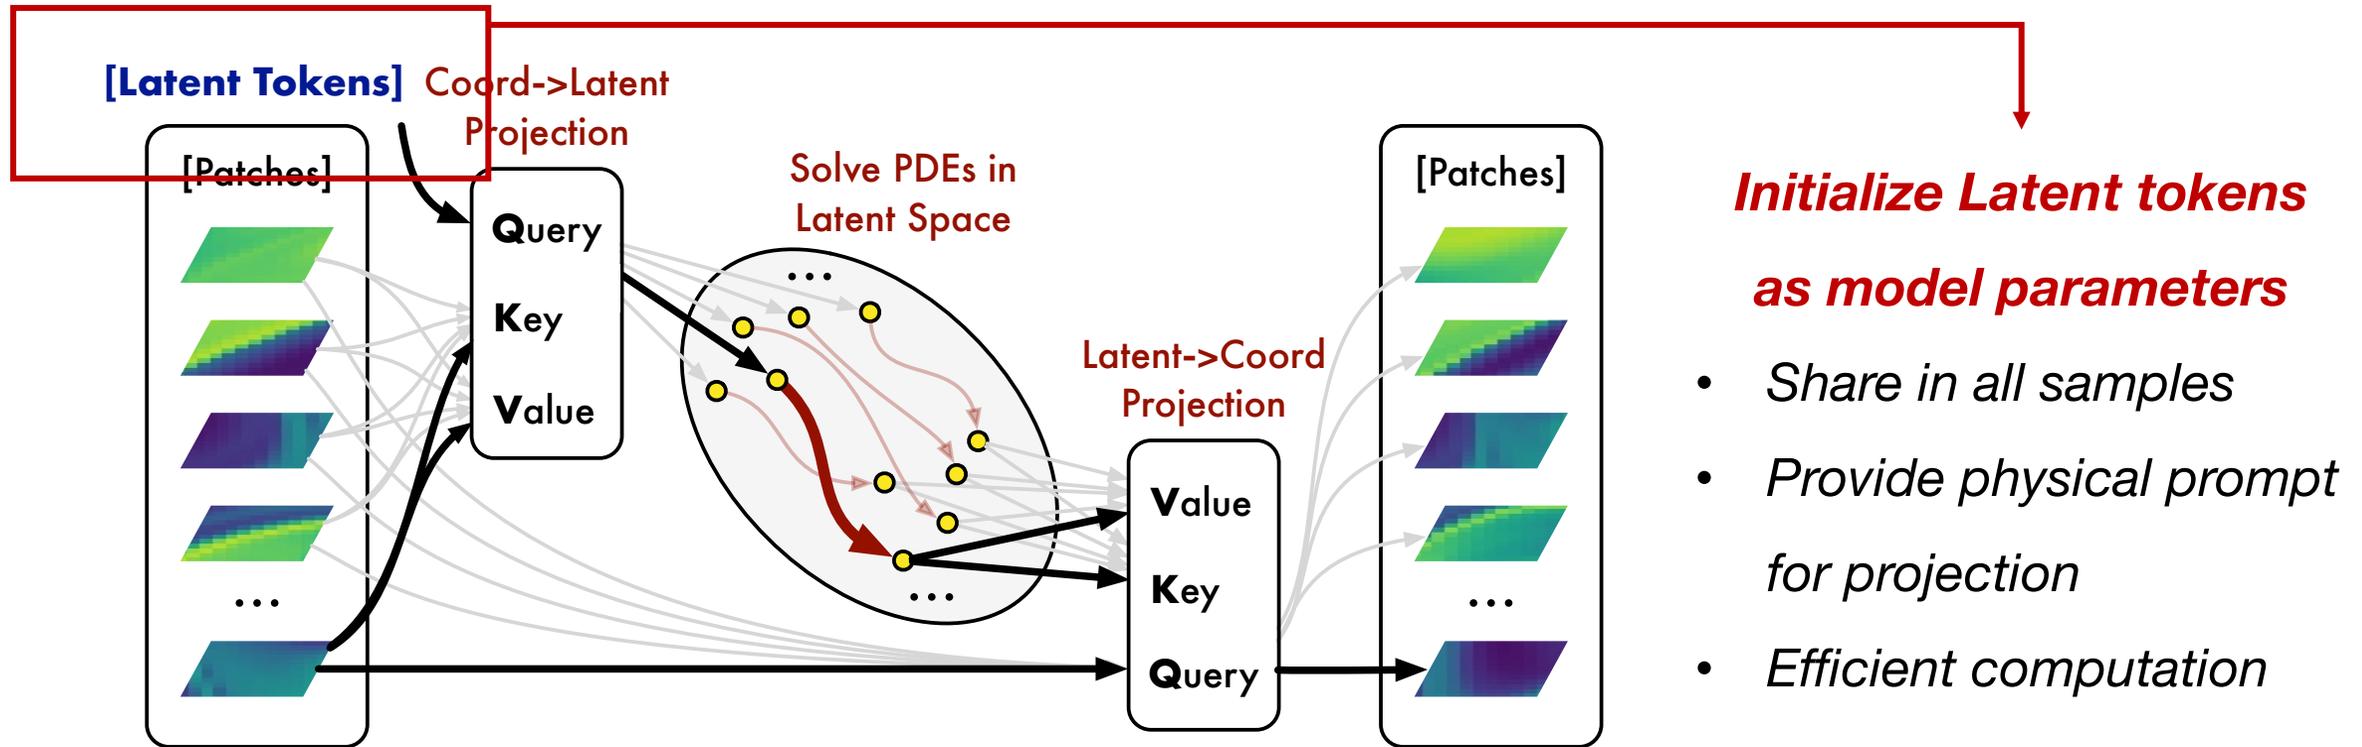
Hierarchical Projection Network



***Still in the coordinate space
Undergoing the problems from
high-dimensional PDEs***

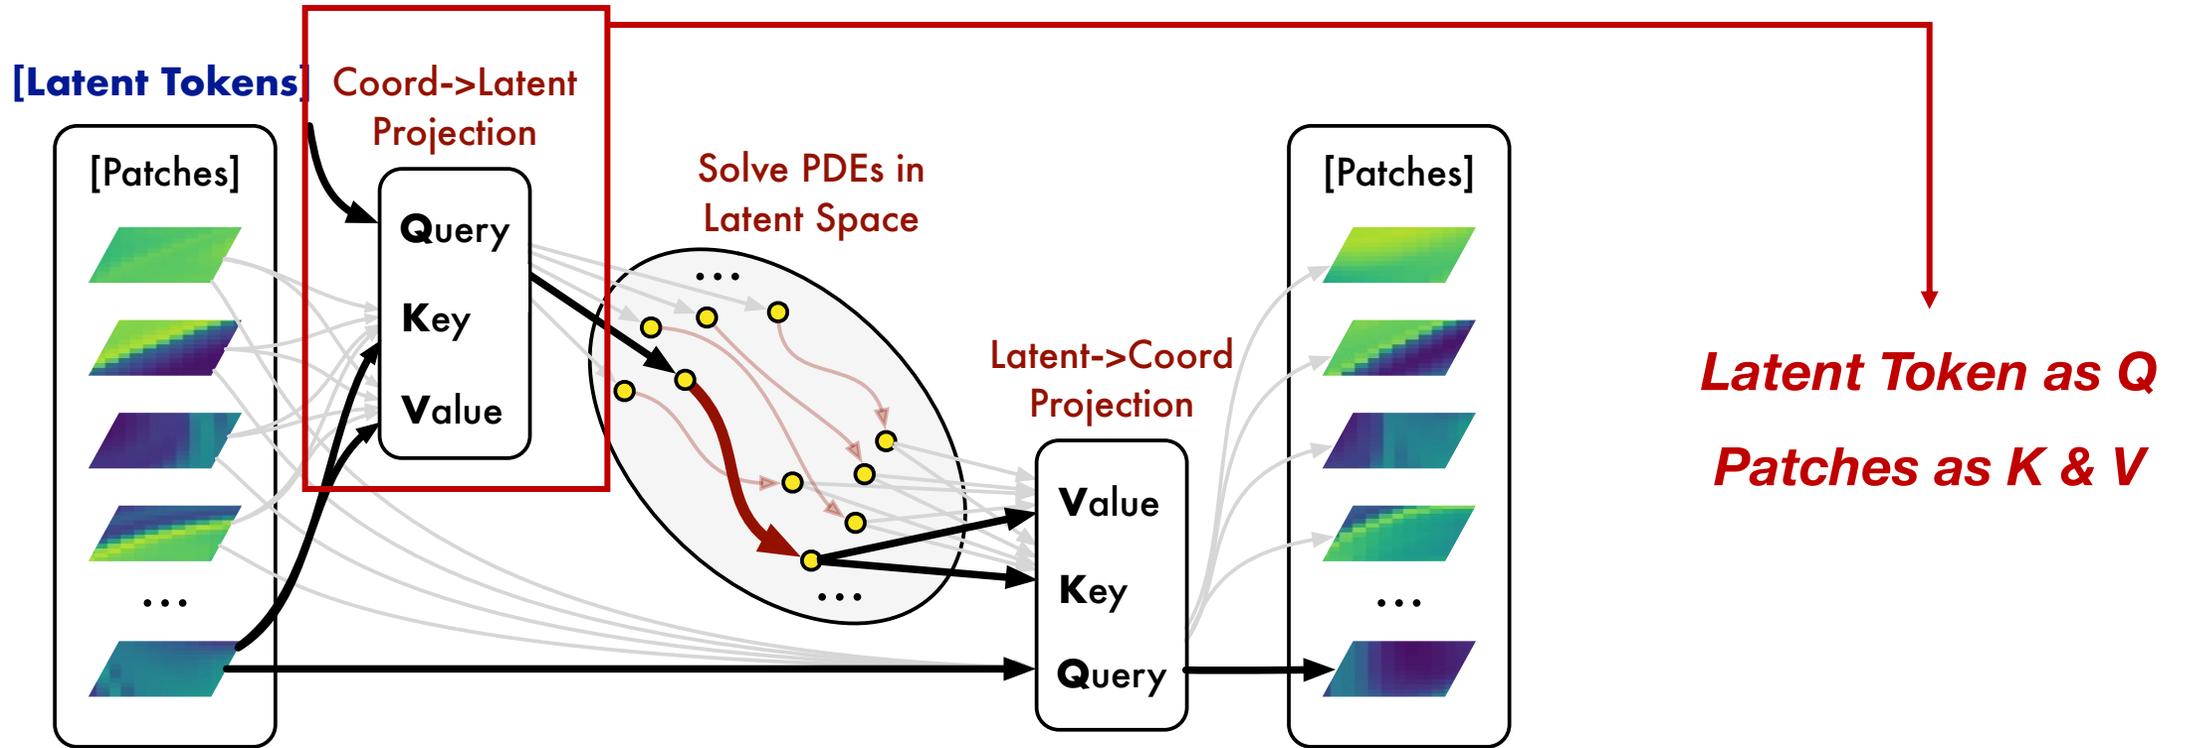
② Attention-based projector → Remove unwieldy coordinate information

Hierarchical Projection Network



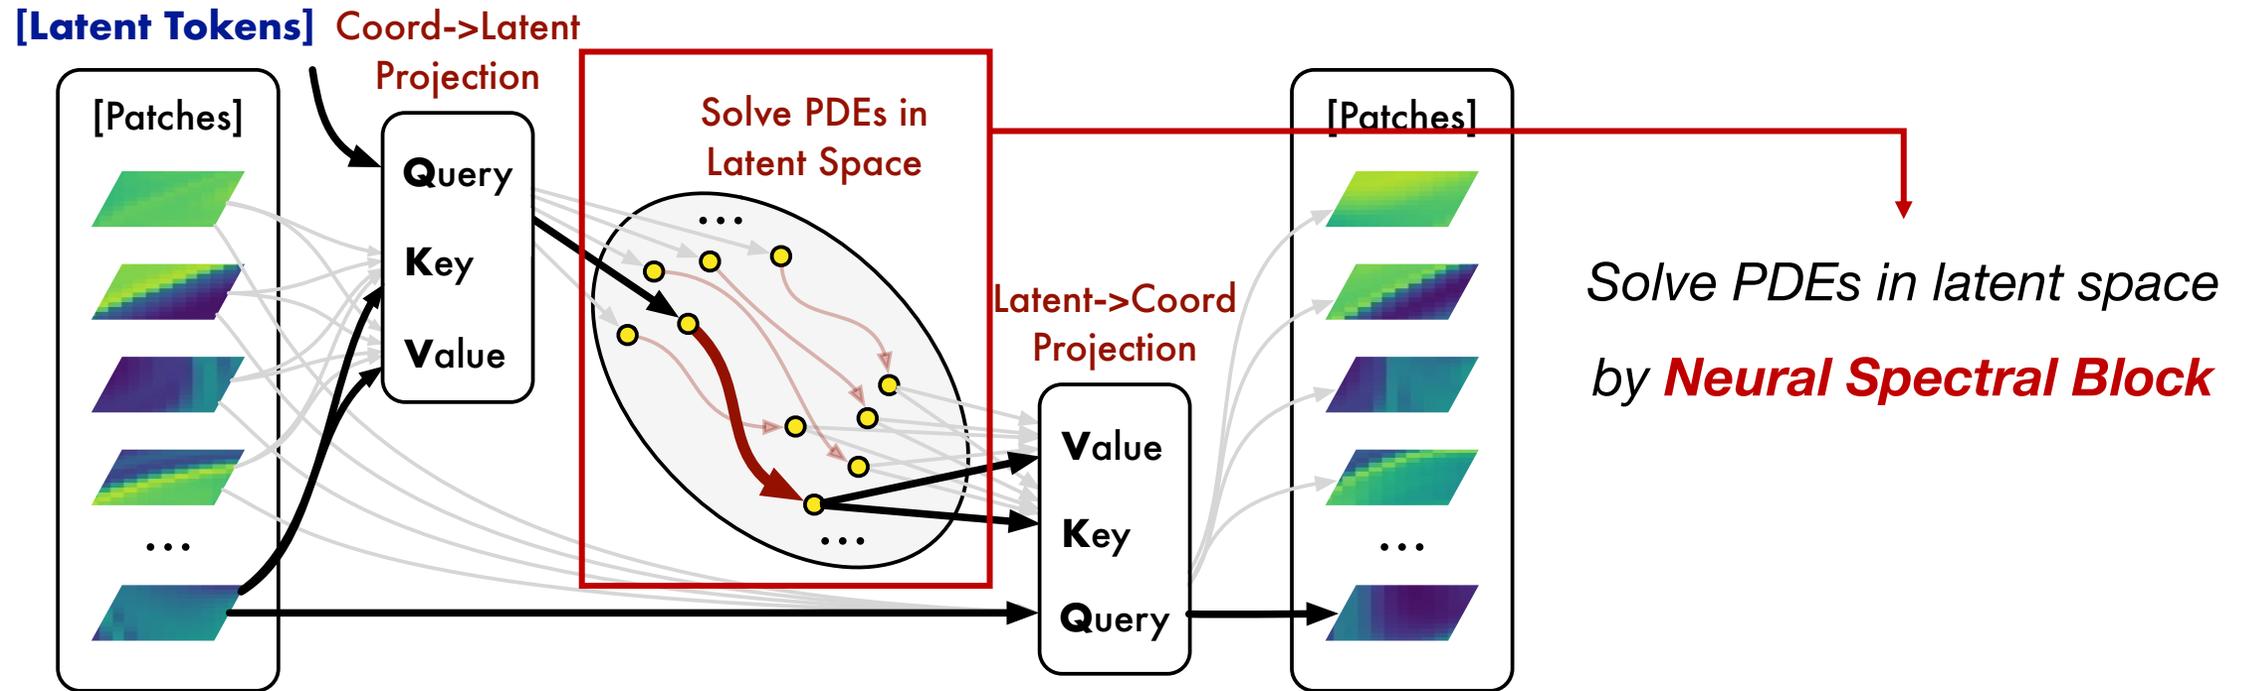
② **Attention-based projector** → *Remove unwieldy coordinate information*

Hierarchical Projection Network



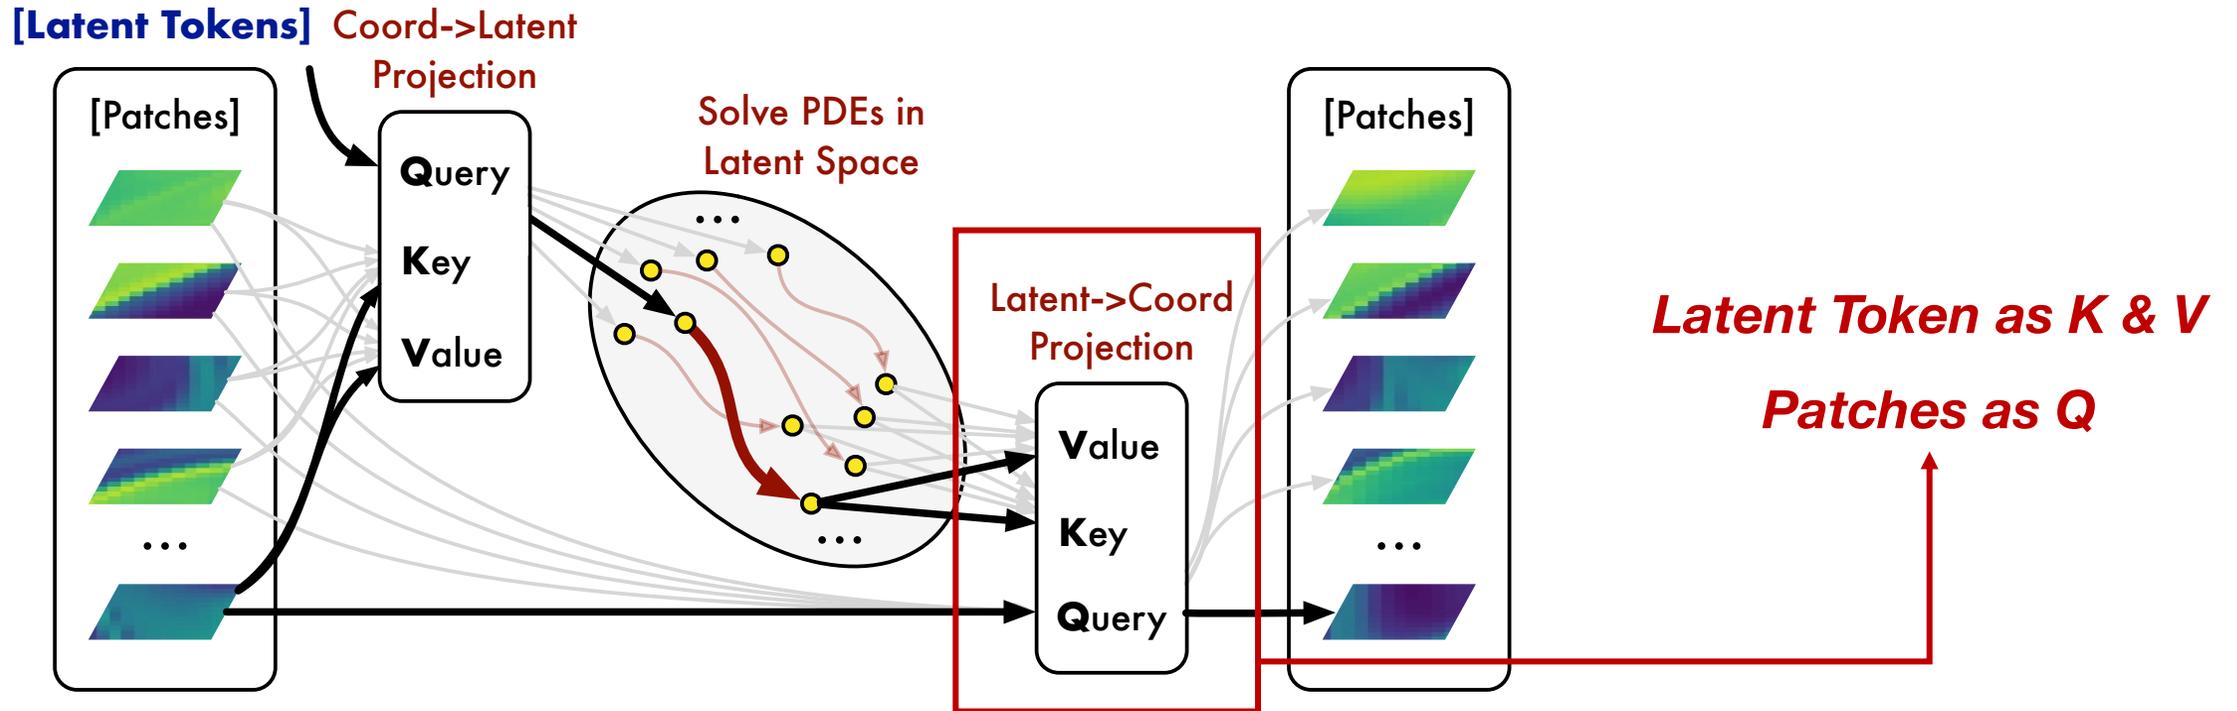
$$\mathbf{T}_{\mathbf{x},i} = \mathbf{T}_i + \sum_{\mathbf{s} \in \mathcal{D}} \frac{\text{Sim}(\mathbf{T}_i, \mathbf{x}(\mathbf{s})\mathbf{W}_K)}{\sum_{\mathbf{s}' \in \mathcal{D}} \text{Sim}(\mathbf{T}_i, \mathbf{x}(\mathbf{s}')\mathbf{W}_K)} (\mathbf{x}(\mathbf{s})\mathbf{W}_V)$$

Hierarchical Projection Network



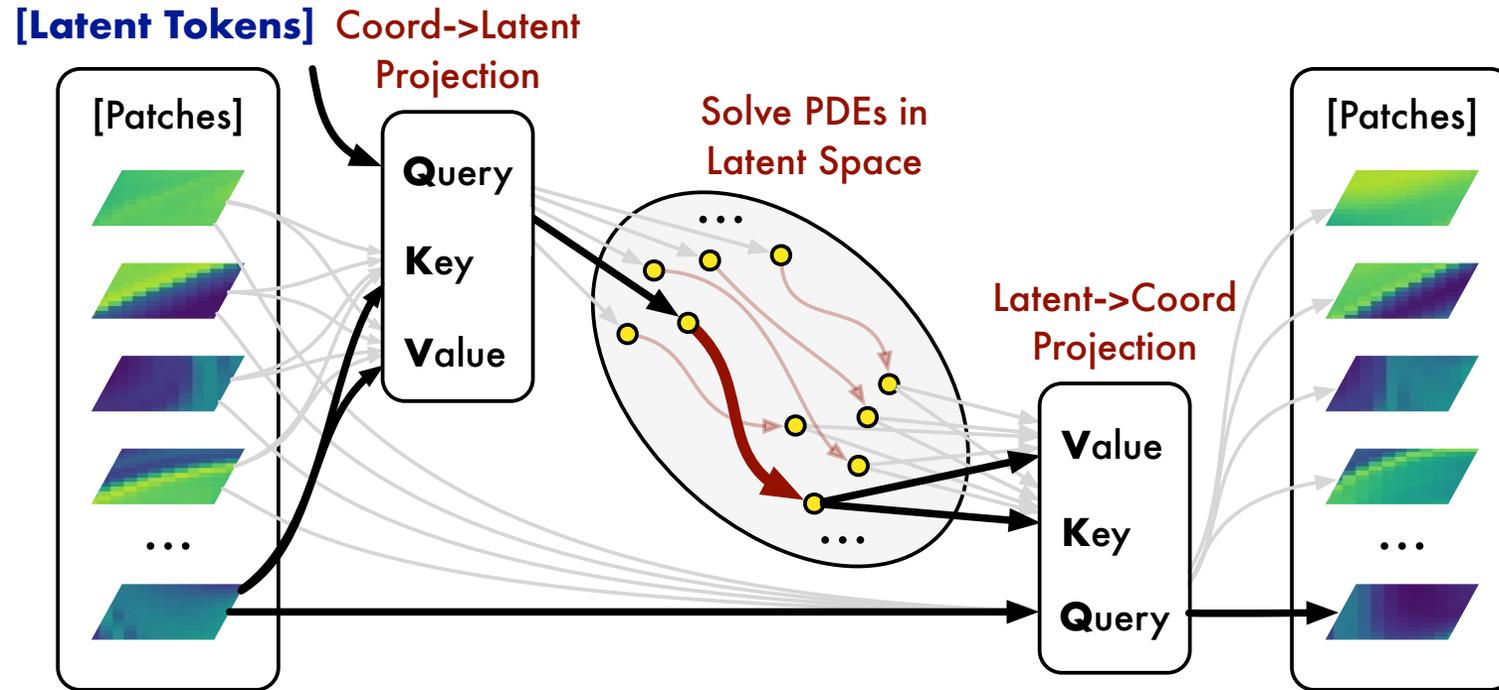
$$\{\mathbf{T}_{\mathbf{y},i,j}^k\}_{i=1}^C = \text{Solve}(\{\mathbf{T}_{\mathbf{x},i,j}^k\}_{i=1}^C)$$

Hierarchical Projection Network



$$\hat{\mathbf{y}}(\mathbf{s}) = \mathbf{x}(\mathbf{s}) + \sum_{i=1}^C \frac{\text{Sim}(\mathbf{x}(\mathbf{s}), \mathbf{T}_{\mathbf{y},i} \mathbf{W}'_{\mathbf{K}})}{\sum_{i'=1}^C \text{Sim}(\mathbf{x}(\mathbf{s}), \mathbf{T}_{\mathbf{y},i'} \mathbf{W}'_{\mathbf{K}})} (\mathbf{T}_{\mathbf{y},i} \mathbf{W}'_{\mathbf{V}})$$

Hierarchical Projection Network



- 1. Linear complexity projection, more efficient computation**
- 2. Highlight the inherent properties of high-dimensional data**
- 3. Benefit the model convergence properties**

Neural Spectral Block

Task: Approximate Complex Nonlinear Mapping



(a) U-Net: Directly Learn Mapping



(b) FNO: Linear Transformation in Fourier Domain

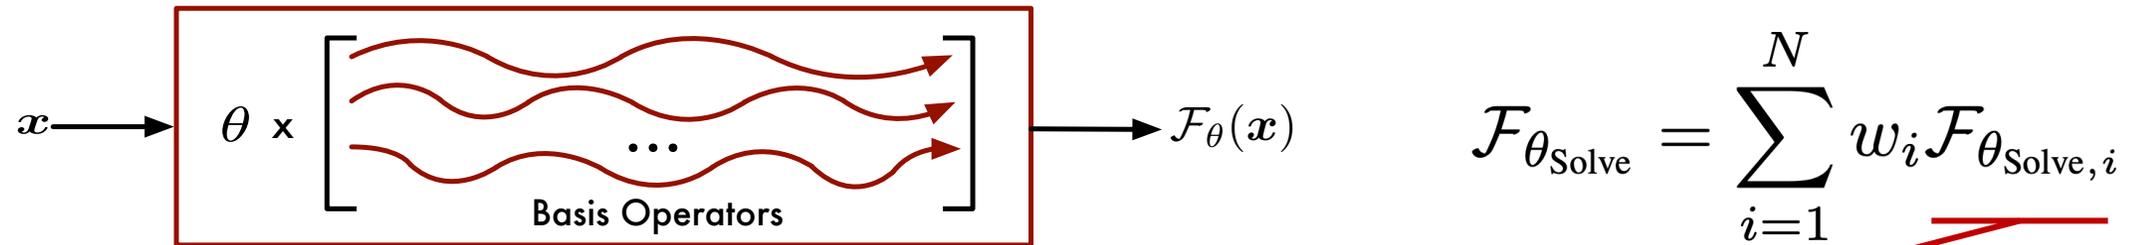


(c) LSM: Decompose into Basis Operators



LSM approximates complex mappings by **learning multiple basis operators**

Neural Spectral Block

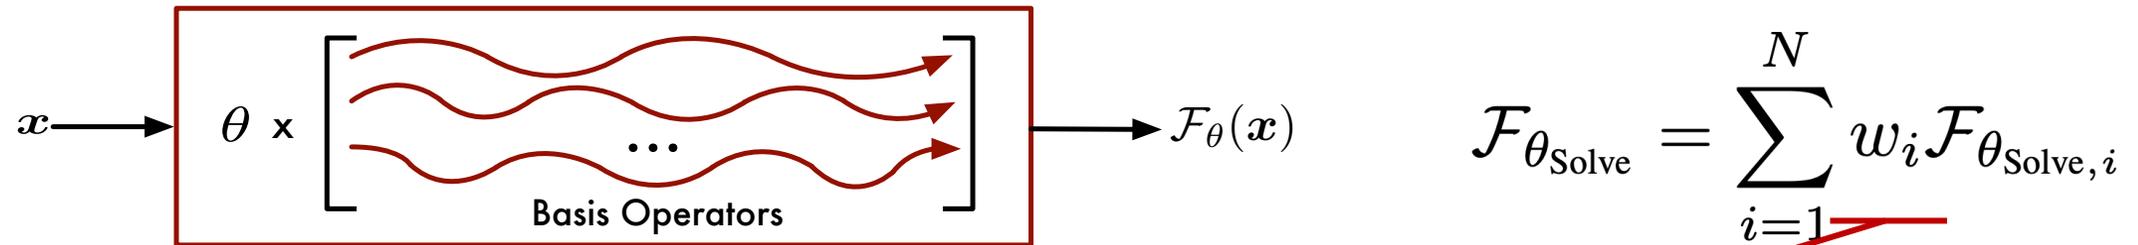


We select trigonometric basis operators

$$\mathcal{F}_{\theta_{\text{Solve}},(2k-1)}(\mathbf{t}_x(\mathbf{s})) = \sin(k\mathbf{t}_x(\mathbf{s}))$$

$$\mathcal{F}_{\theta_{\text{Solve}},(2k)}(\mathbf{t}_x(\mathbf{s})) = \cos(k\mathbf{t}_x(\mathbf{s}))$$

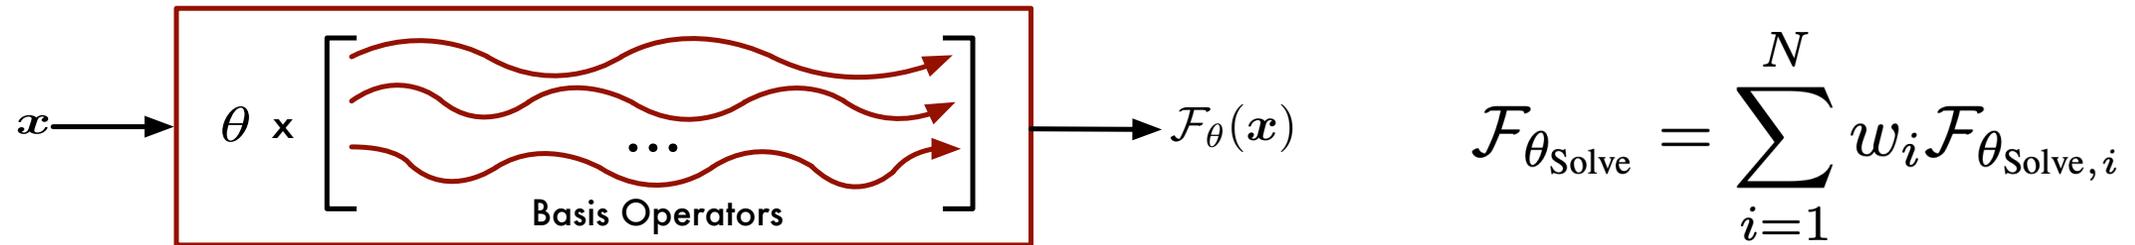
Neural Spectral Block



$\{\omega_1, \dots, \omega_N\}$ are model parameters

Will be optimized to satisfy the PDEs better,
during the training process, namely solving PDEs

Neural Spectral Block



Neural spectral block is applied to
projected latent tokens of **all the patches in multiple scales**

$$\mathbf{T}_y = \mathbf{T}_x + \mathbf{w}_0 + \mathbf{w}_{\sin} \begin{bmatrix} \sin(\mathbf{T}_x) \\ \vdots \\ \sin(\frac{N}{2} \mathbf{T}_x) \end{bmatrix} + \mathbf{w}_{\cos} \begin{bmatrix} \cos(\mathbf{T}_x) \\ \vdots \\ \cos(\frac{N}{2} \mathbf{T}_x) \end{bmatrix}$$

Theoretical analysis

Convergence of Trigonometric Approximation in High-dimensional Space:

Let $f: \mathbb{R}^M \rightarrow \mathbb{R}^M$ be a 2π -periodic function w.r.t. the variable on each dimension, where $f \in L_p([-\pi, \pi)^M)$, $M \geq 2$, $1 \leq p \leq \infty$ and $p \neq 2$. For f defined on the **M -dimension space**, its trigonometric approximation f^N is defined as:

$$f^N(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^M, |\mathbf{k}| \leq N} \left(\frac{1}{2\pi} \int_{[-\pi, \pi)^M} f(\mathbf{t}) e^{-i\mathbf{k}\mathbf{t}} d\mathbf{t} \right) e^{i\mathbf{k}\mathbf{x}},$$

if f satisfies the Lipschitz condition, then there exists a constant K , such that

$$\|f - f^N\| \leq KN^{(M-1)\left|\frac{1}{2}-\frac{1}{p}\right|-1}.$$

Slow Convergence Rate in High-dimensional Space

Theoretical analysis

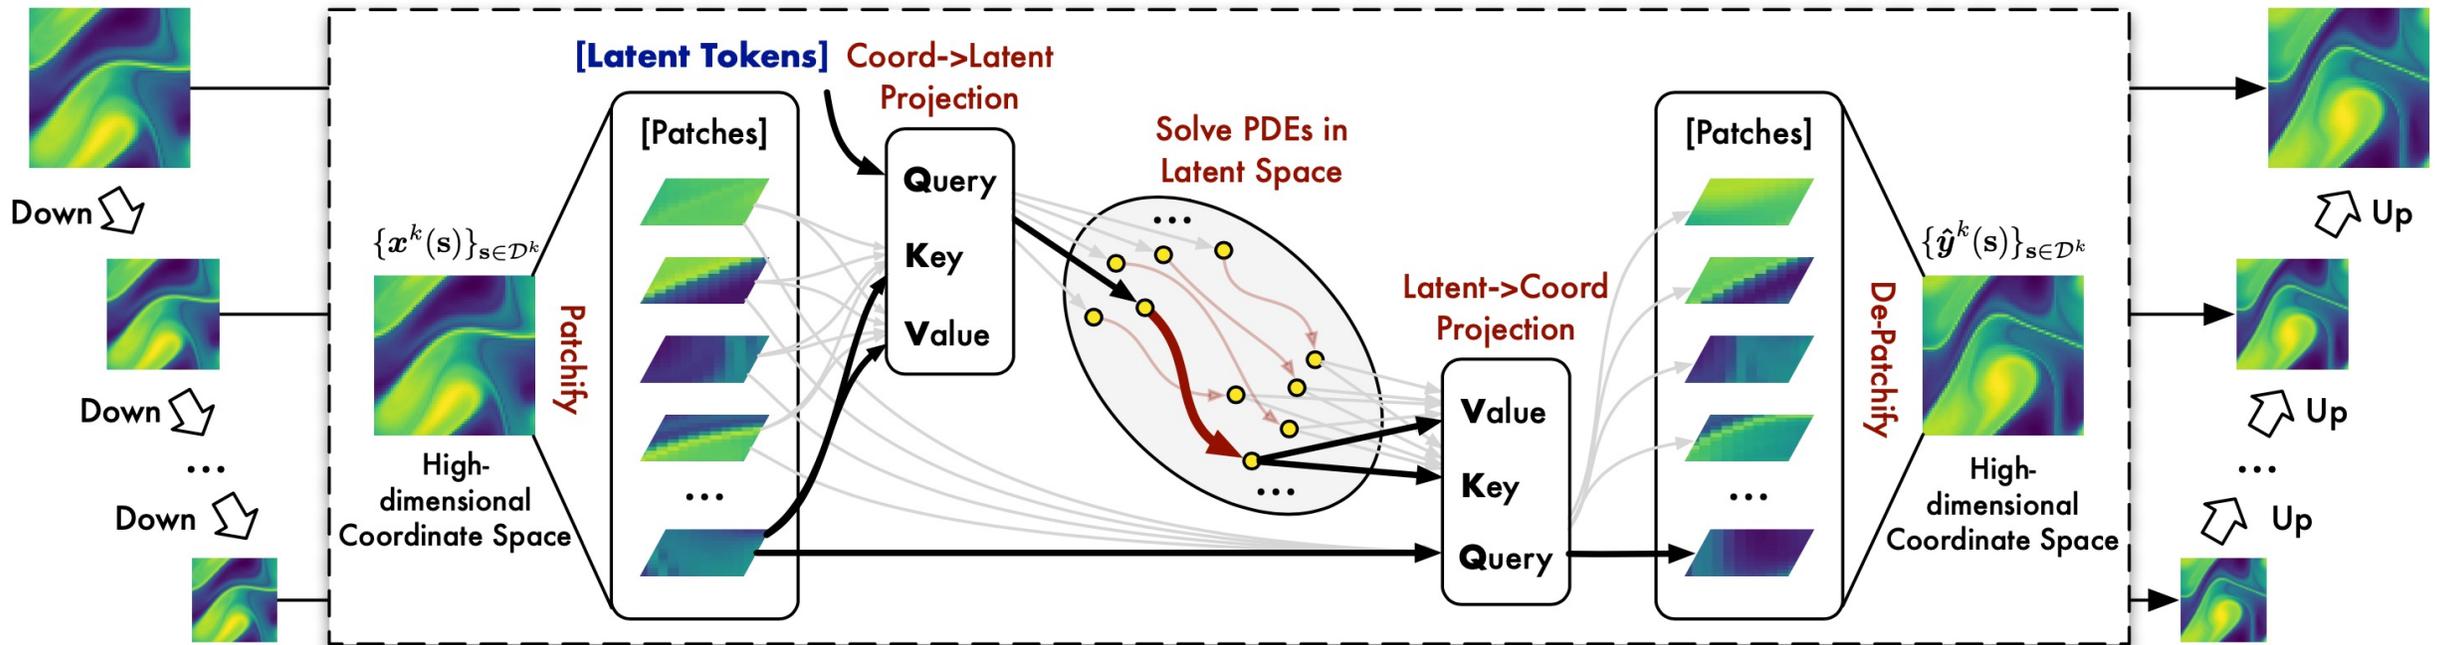
Approximation and Convergence Properties of *Neural Spectral Block*

(trigonometric approximation with residual): Given $f: [0, \pi] \rightarrow \mathbb{R}$, if f satisfies the Lipschitz condition, there is a choice of model parameters such that the approximation f^N defined in neural spectral block can uniformly converge to f with the speed as follows:

$$\|f - f^N(\mathbf{x})\| \leq K \frac{\ln N}{N}, \forall x \in [0, \pi].$$

Projecting M-dimension data into **independent latent tokens** brings favorable convergence speed of Neural Spectral Block.

Overall design of LSM



LSM with *Hierarchical Projection Network* and *Neural Spectral Block*

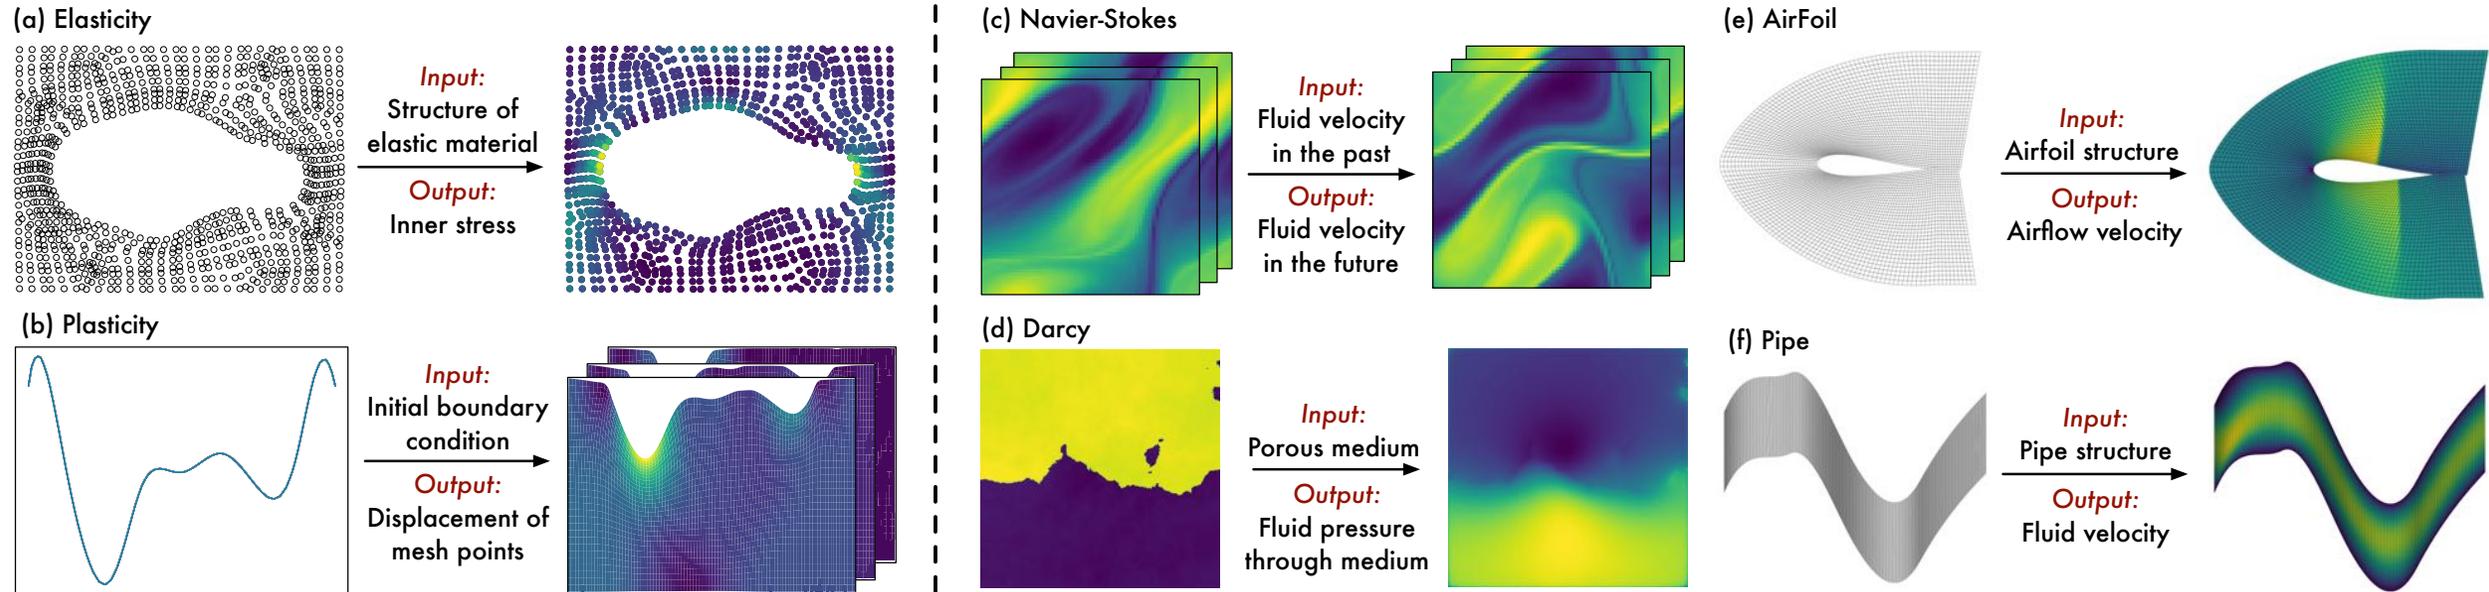
① Coor \rightarrow Latent ② Solving in the Latent Space ③ Latent \rightarrow Coor

Experiments

PHYSICS	BENCHMARKS	GEOMETRY	#DIM
SOLID	ELASTICITY-P	POINT CLOUD	2D
	ELASTICITY-G	REGULAR GRID	2D
	PLASTICITY	STRUCTURED MESH	3D
FLUID	NAVIER-STOKES	REGULAR GRID	3D
	DARCY	REGULAR GRID	2D
	AIRFOIL	STRUCTURED MESH	2D
	PIPE	STRUCTURED MESH	2D

Seven typical PDE solving tasks, covering both fluid and solid physics, various geometrics.

PDE-governed Tasks



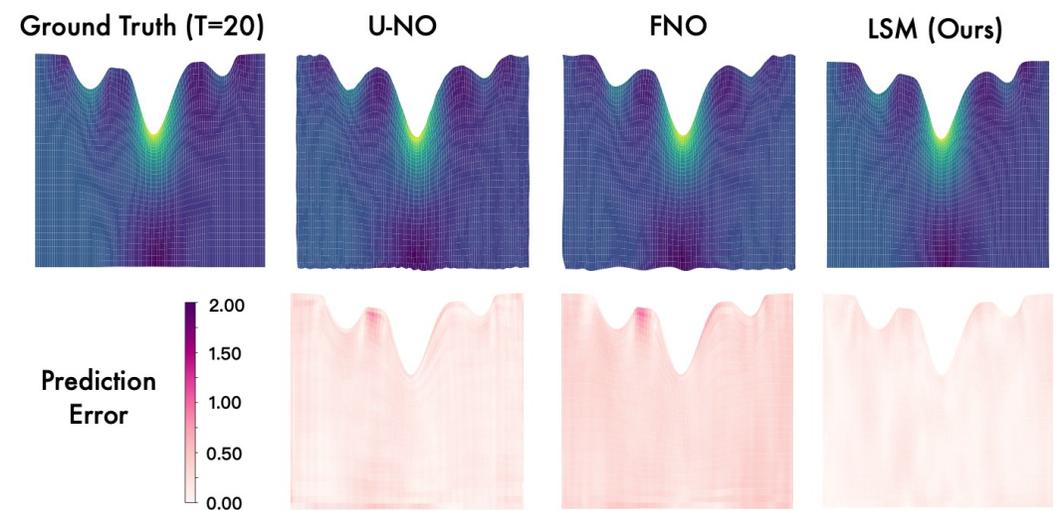
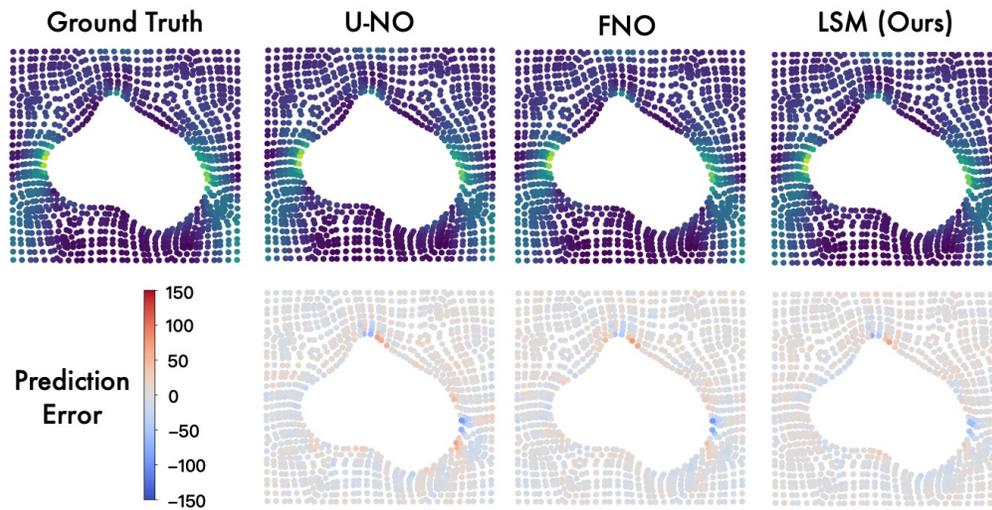
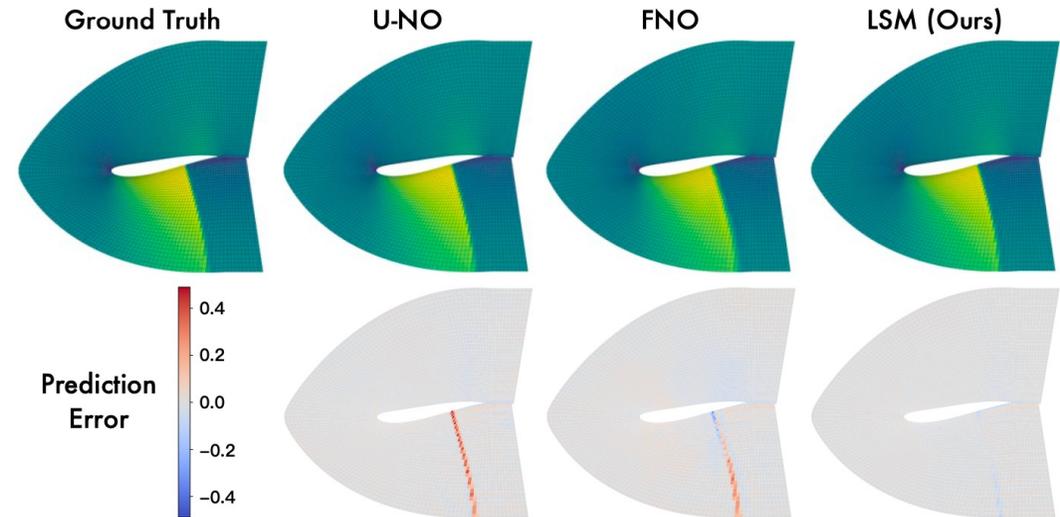
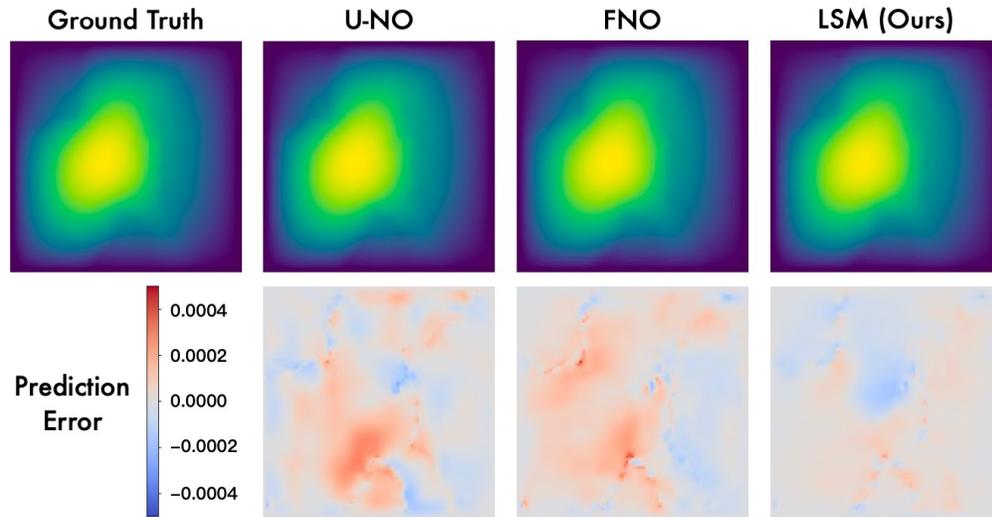
Approximate complex input-output mappings with deep models

Main Results

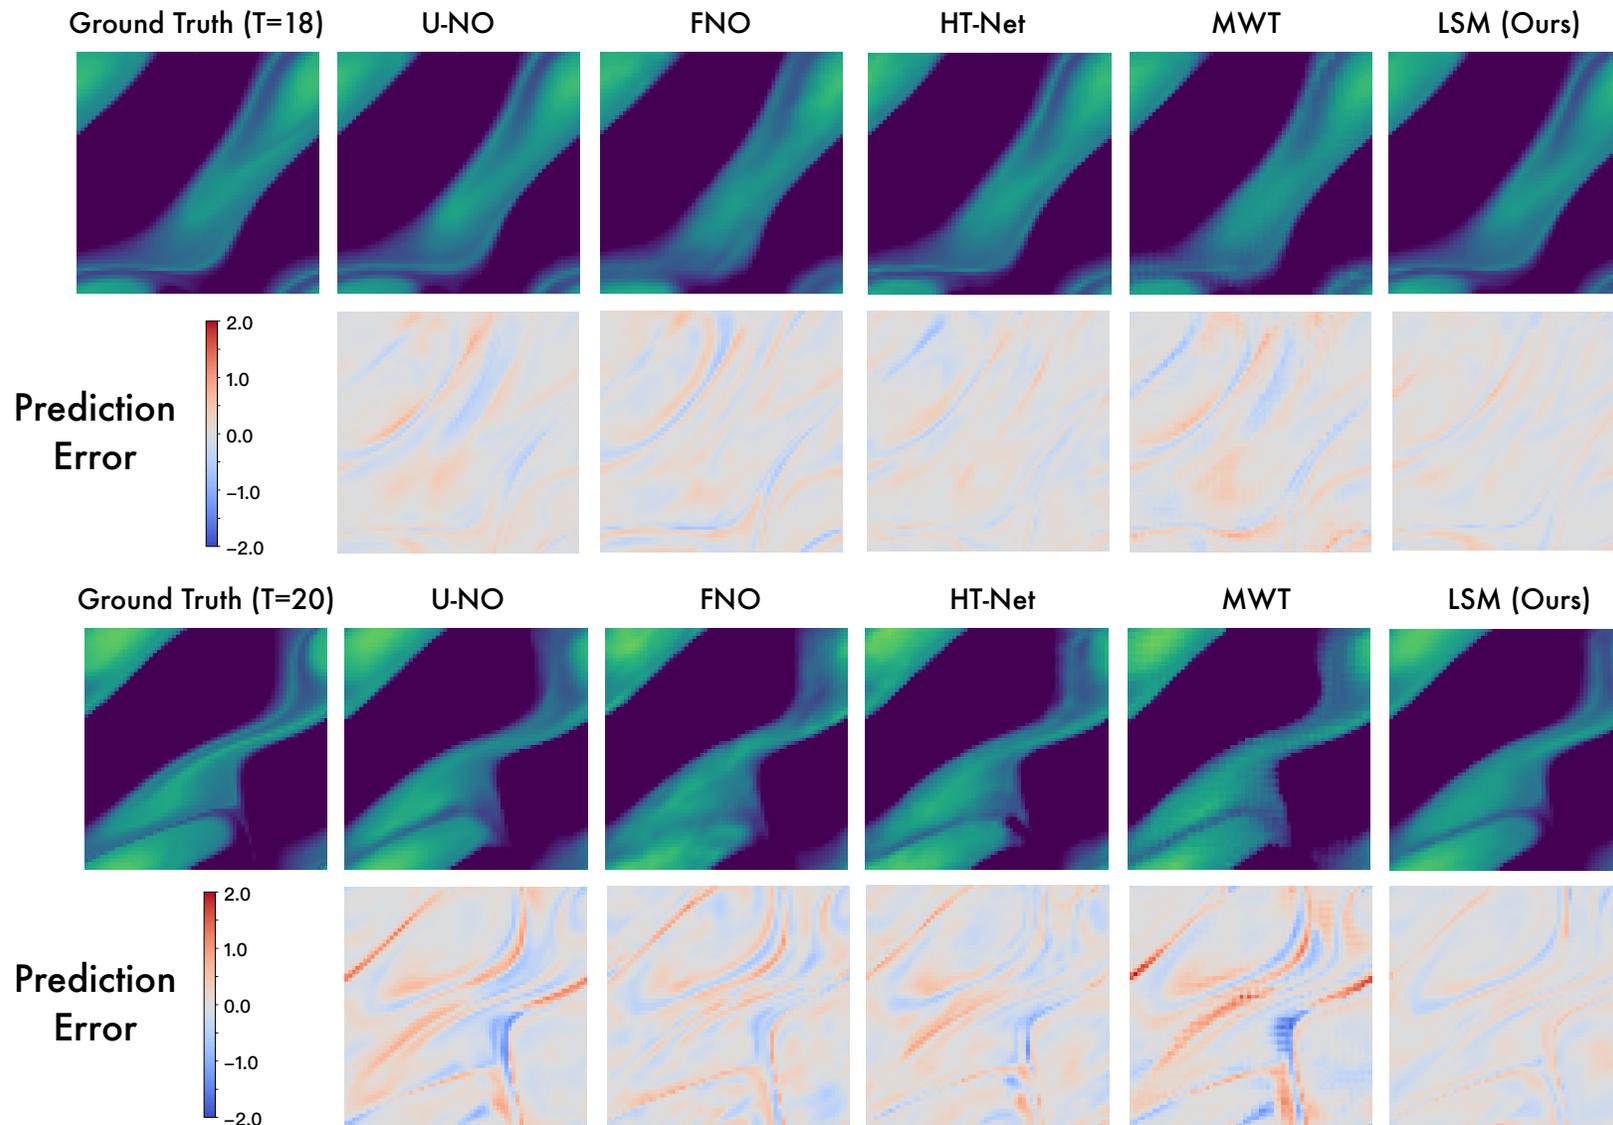
MODEL	SOLID PHYSICS*			FLUID PHYSICS†			
	ELASTICITY-P ‡	ELASTICITY-G	PLASTICITY	NAVIER–STOKES	DARCY	AIRFOIL	PIPE
U-NET (2015)	0.0235	0.0531	0.0051	0.1982	0.0080	0.0079	0.0065
RESNET (2016)	0.0262	0.0843	0.0233	0.2753	0.0587	0.0391	0.0120
TF-NET (2019)	/	/	/	0.1801	/	/	/
SWIN (2021)	0.0283	0.0819	0.0170	0.2248	0.0397	0.0270	0.0109
DEEPONET (2021)	0.0965	0.0900	0.0135	0.2972	0.0588	0.0385	0.0097
FNO (2021)	<u>0.0229</u>	0.0508	0.0074	0.1556	0.0108	0.0138	0.0067
U-FNO (2021)	0.0239	0.0480	0.0039	0.2231	0.0183	0.0269	<u>0.0056</u>
WMT (2021)	0.0359	0.0520	0.0076	<u>0.1541</u>	0.0082	0.0075	<u>0.0077</u>
GALERKIN (2021)	0.0240	0.1681	0.0120	0.2684	0.0170	0.0118	0.0098
SNO (2022)	0.0390	0.0987	0.0070	0.2568	0.0495	0.0893	0.0294
U-NO (2022)	0.0258	<u>0.0469</u>	<u>0.0034</u>	0.1713	0.0113	0.0078	0.0100
HT-NET (2022)	0.0372	0.0472	0.0333	0.1847	0.0079	<u>0.0065</u>	0.0059
F-FNO (2023)	0.0263	0.0475	0.0047	0.2322	<u>0.0077</u>	0.0078	0.0070
KNO (2023A)	/	/	/	0.2023	/	/	/
LSM	0.0218	0.0408	0.0025	0.1535	0.0065	0.0059	0.0050
PROMOTION	4.8%	13.0%	26.5%	0.4%	15.6%	9.2%	10.7%

**LSM achieves consistent SOTA and surpasses previous
14 baselines with 11.5% error reduction.**

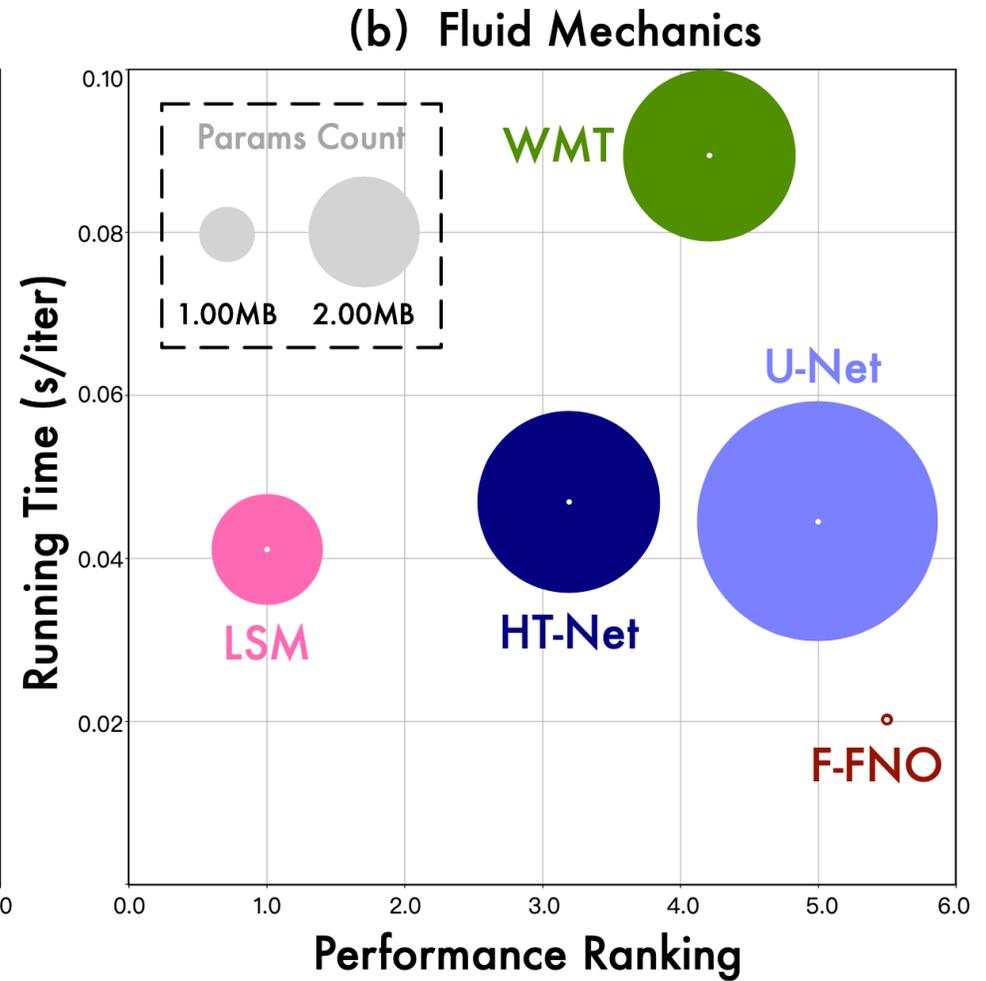
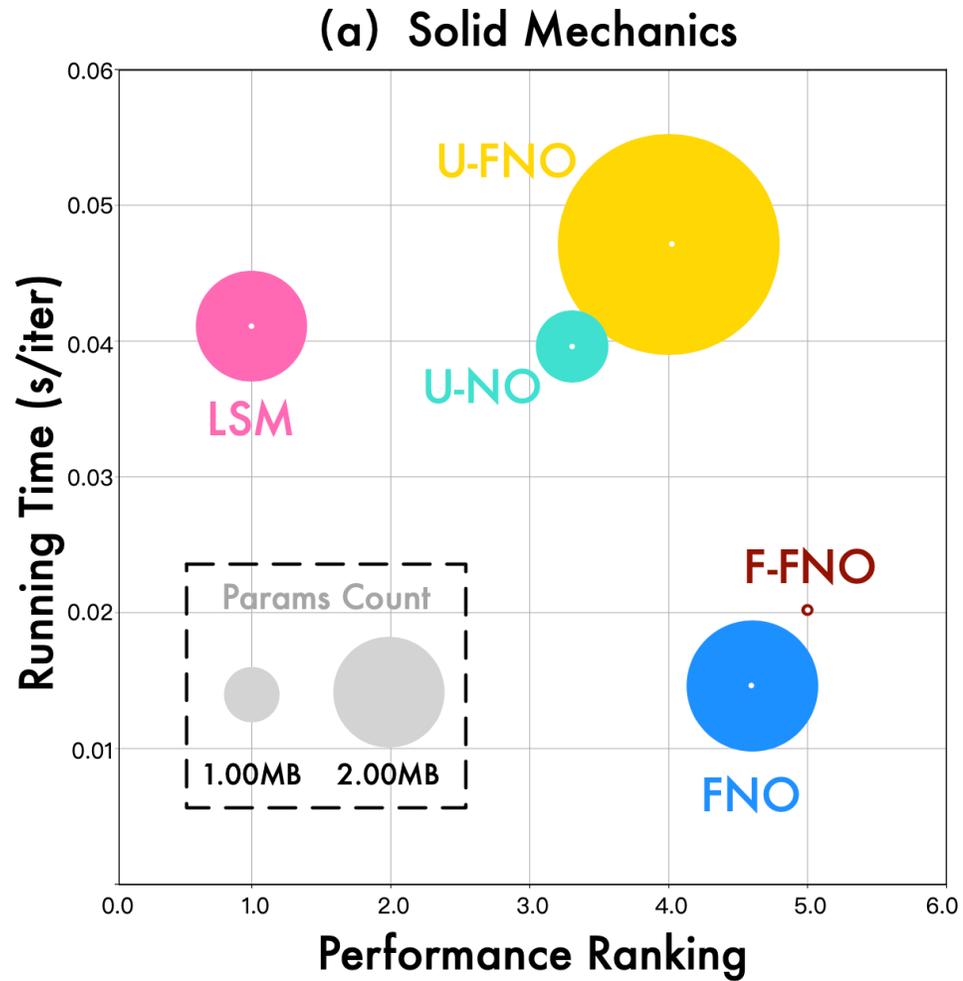
Showcases



Showcases

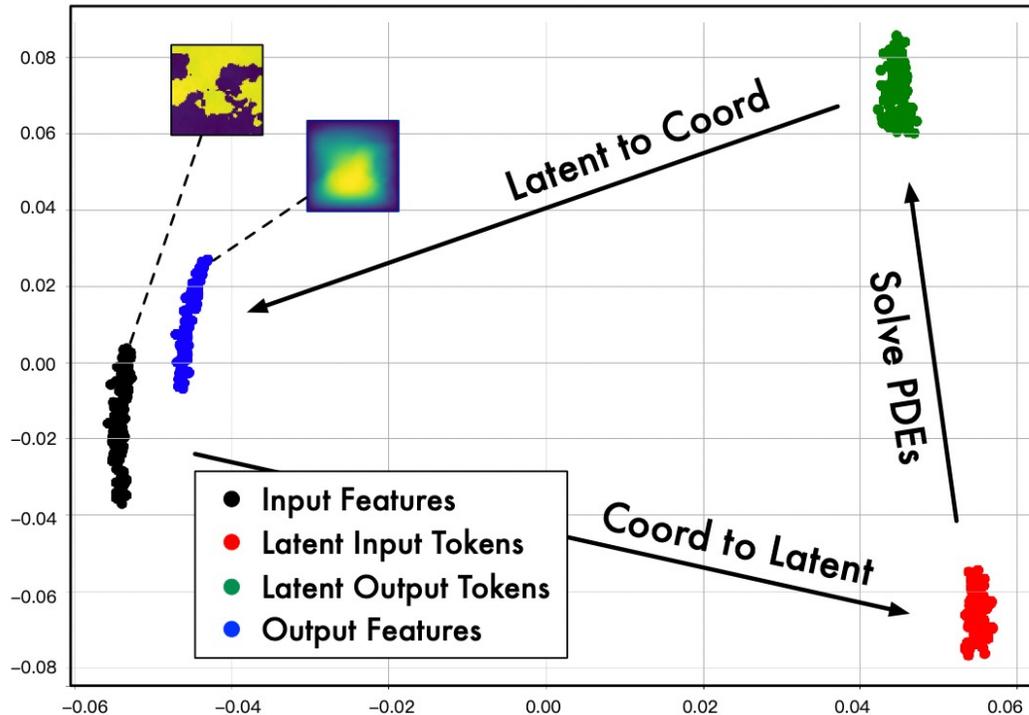


Efficiency

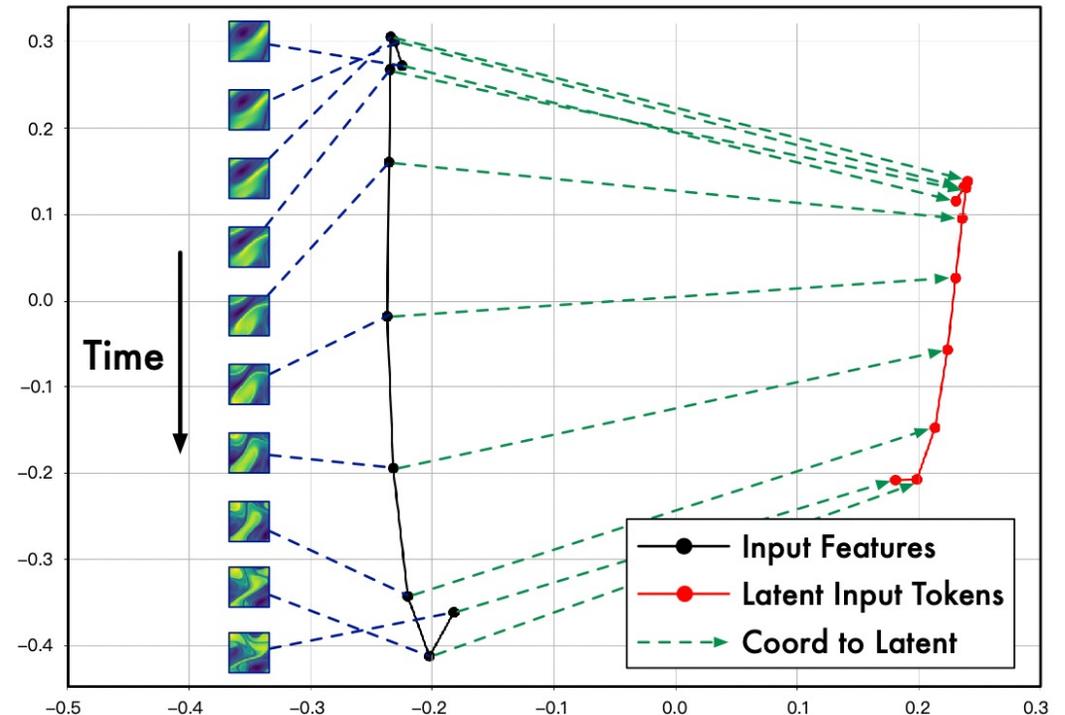


Favorable trade-off between performance and efficiency.

Solving Process Visualization



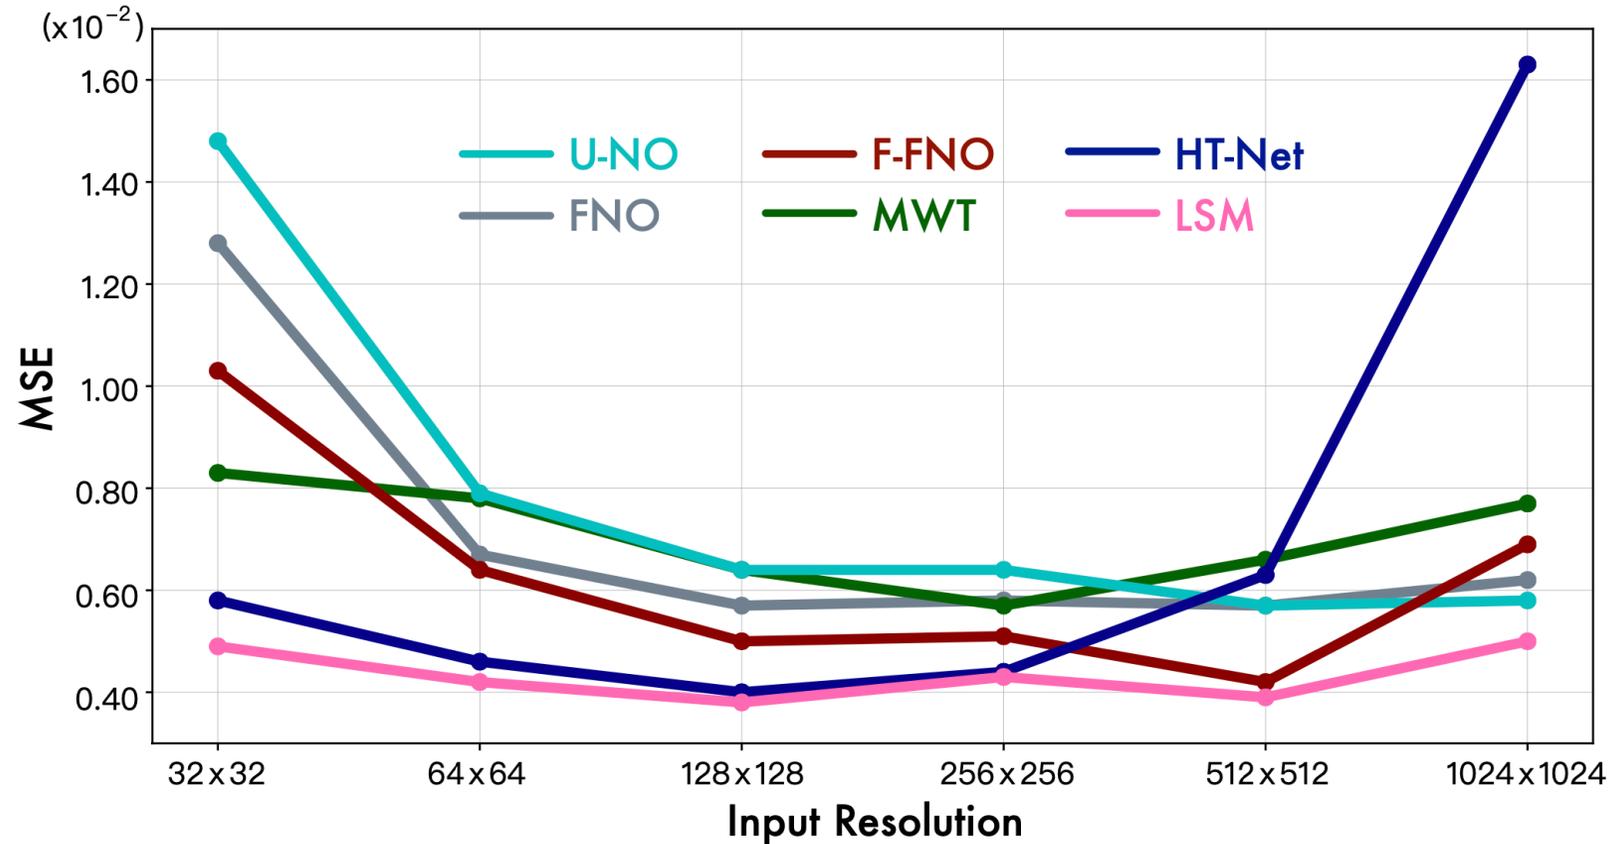
(a) Darcy Benchmark



(b) Latent Process on Navier-Stokes

LSM can precisely capture the complex mapping and latent process from high-dimensional coordinate space.

Performance under various resolutions in Darcy benchmark

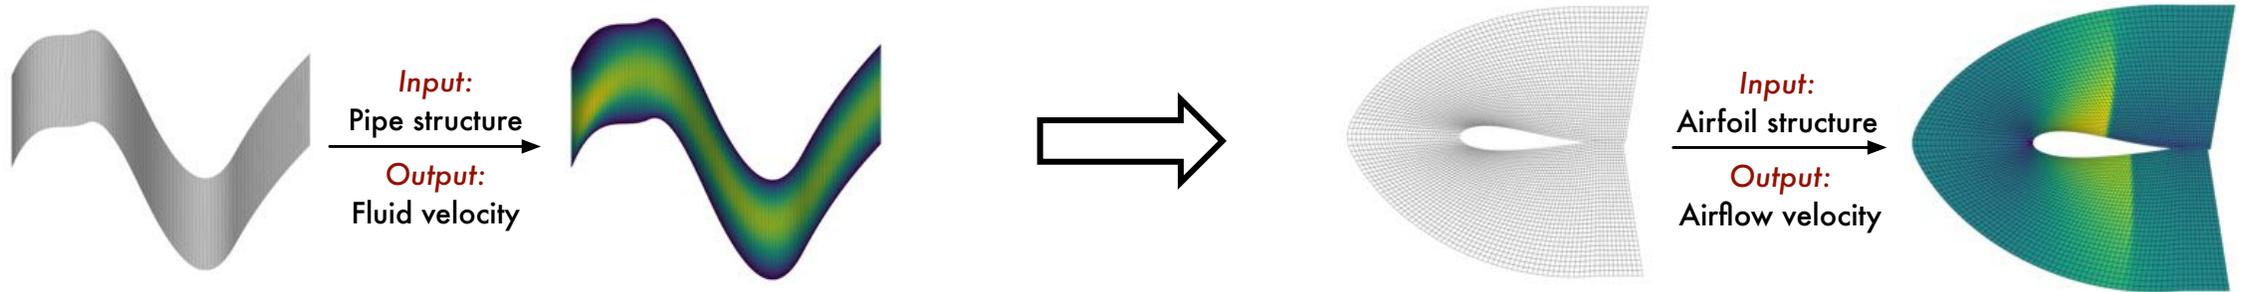


LSM presents a stable performance w.r.t. different inputs and consistently surpasses other baselines.

Transferability

Finetune the **pipe-pretrained** model into **airfoil with limited data**

(same PDE equation but different boundary conditions)



MSE ($\times 10^{-2}$)	20% AIRFOIL DATA	40% AIRFOIL DATA	60% AIRFOIL DATA	80% AIRFOIL DATA	100% AIRFOIL DATA
U-NET (2015)	1.88 → 1.93 (-2.7%)	1.38 → 1.14 (+17.3%)	0.96 → 0.90 (+6.3%)	0.85 → 0.81 (+4.7%)	0.79 → 0.77 (+2.5%)
U-NO (2022)	6.30 → 1.72	2.39 → 1.73	1.10 → 1.00 (+9.1%)	0.86 → 0.82 (+4.7%)	0.78 → 0.82 (-5.1%)
HT-NET (2022)	1.73 → 1.43 (+17.3%)	1.08 → 0.82 (+24.1%)	0.75 → 0.69 (+8.0%)	0.70 → 0.65 (+7.1%)	0.65 → 0.61 (+6.2%)
LSM	1.66 → 1.31 (+21.1%)	0.91 → 0.75 (+17.6%)	0.69 → 0.61 (+11.6%)	0.63 → 0.58 (+7.9%)	0.59 → 0.55 (+6.8%)

LSM shows good Transferability between different conditions.

Open Source

The screenshot shows the GitHub interface for the repository 'thuml / Latent-Spectral-Models'. The repository is public and has 4 watchers, 2 forks, and 21 stars. The main branch is 'main'. The repository contains a README.md file, a LICENSE file, a .gitignore file, and several Python scripts in the 'scripts' directory. The README.md file is open, showing the title 'Latent Spectral Models (ICML 2023)' and a link to the paper. The right sidebar contains information about the repository, including the MIT license, activity, and suggested workflows.

File	Commit Message	Commit Date
fig	add readme	last month
models	clean	last month
scripts	Update darcy_lsm.sh	last month
utils	clean	last month
.gitignore	Initial commit	last month
LICENSE	Initial commit	last month
README.md	Update README.md	last month
exp_airfoils.py	clean	last month
exp_darcy.py	clean	last month
exp_elas.py	clean	last month
exp_elas_interp.py	clean	last month
exp_ns.py	clean	last month
exp_pipe.py	clean	last month
exp_plas.py	clean	last month
model_dict.py	clean	last month
requirements.txt	init	last month

README.md

Latent Spectral Models (ICML 2023)

Solving High-Dimensional PDEs with Latent Spectral Models [\[paper\]](#)

Python 96.1% **Shell** 3.9%

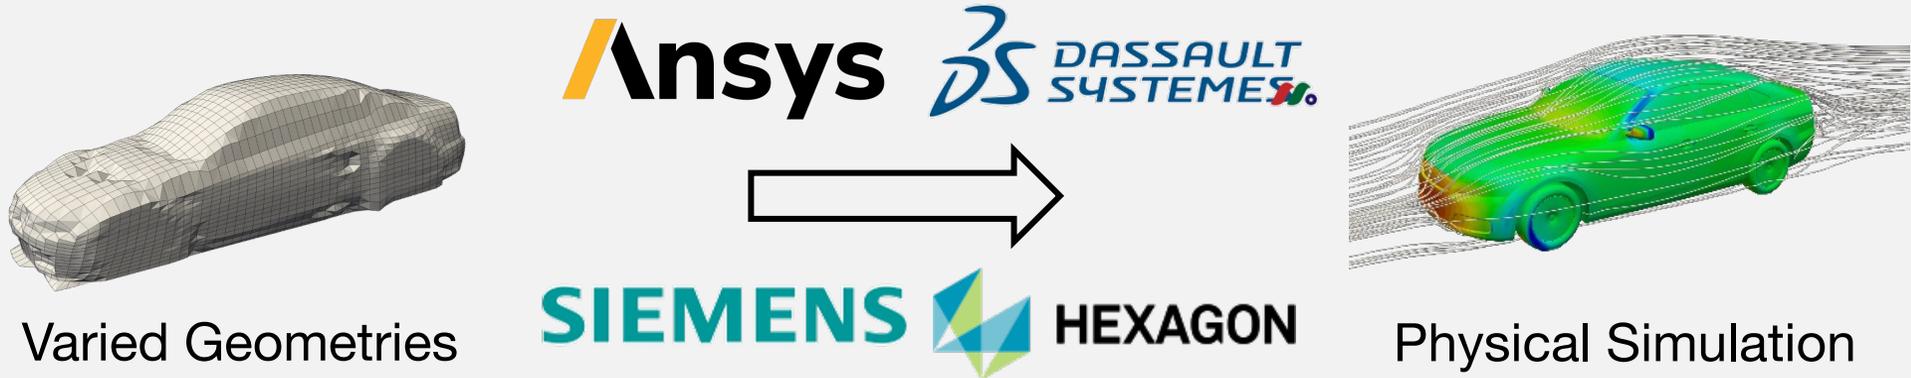
Suggested Workflows
Based on your tech stack

Actions Importer [Set up](#)

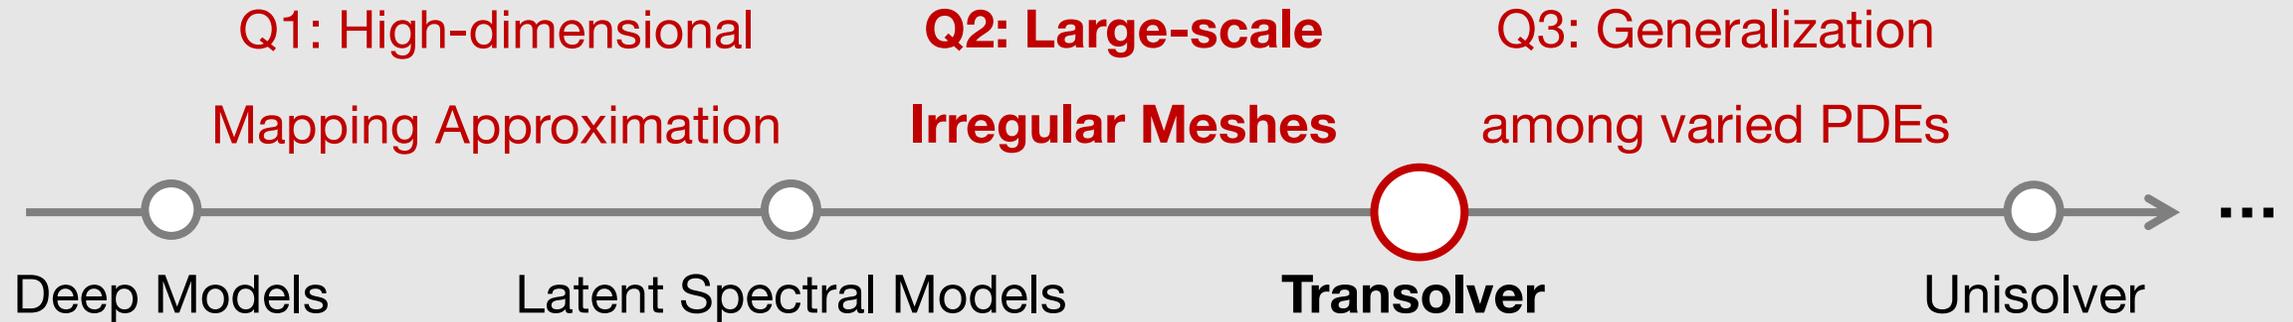
Code is available at <https://github.com/thuml/Latent-Spectral-Models>

A Roadmap to Practical Neural PDE Solvers

Industrial simulation with CAE



Neural PDE Solver (Our work)





ICML | 2024

The Forty-first International Conference on Machine Learning



Transolver: A Fast Transformer Solver for PDEs on General Geometries

Haixu Wu¹ Huakun Luo¹ Haowen Wang¹ Jianmin Wang¹ Mingsheng Long¹



Haixu Wu



Huakun Luo



Haowen Wang

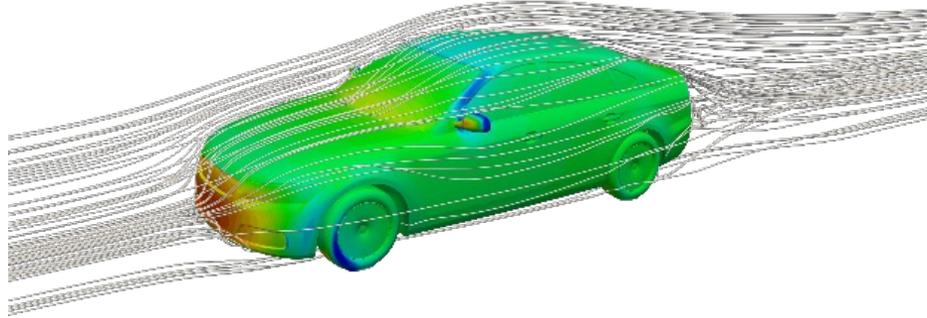


Jianmin Wang

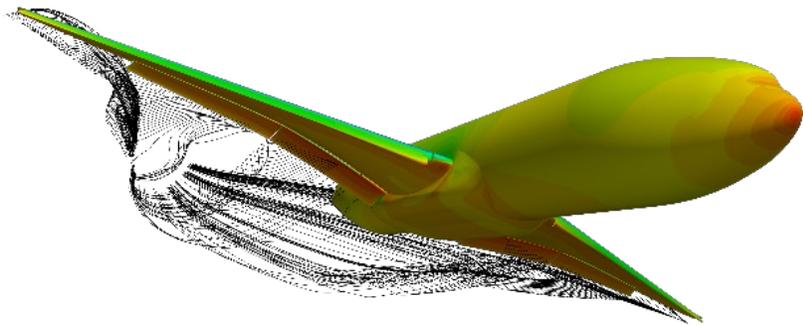
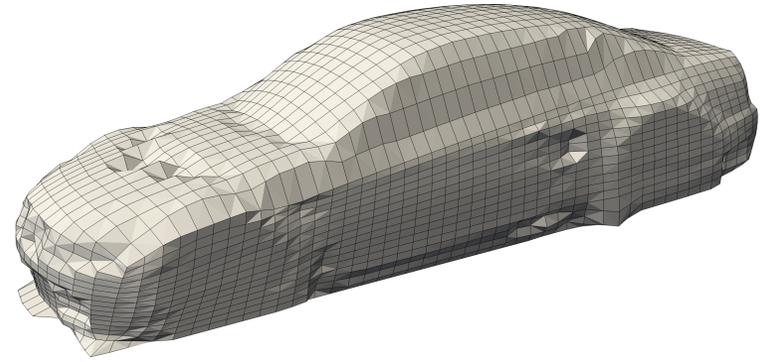
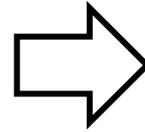


Mingsheng Long

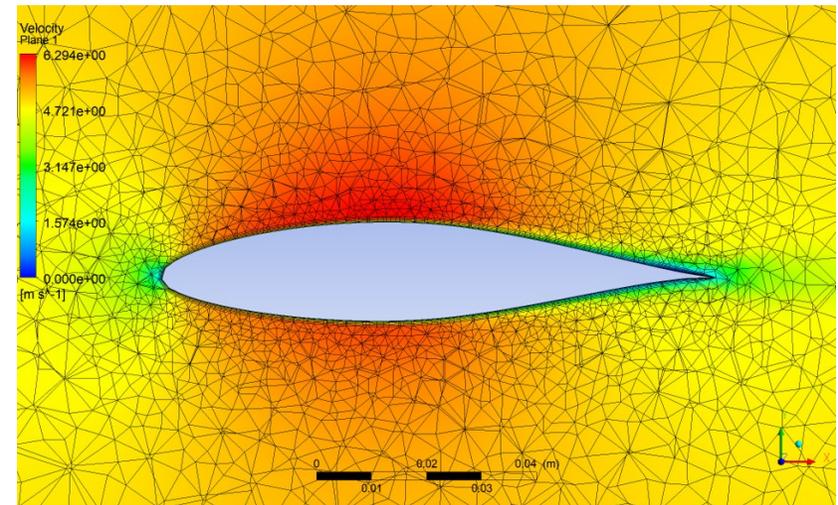
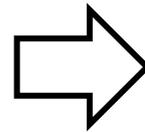
Solving PDEs: Discretization



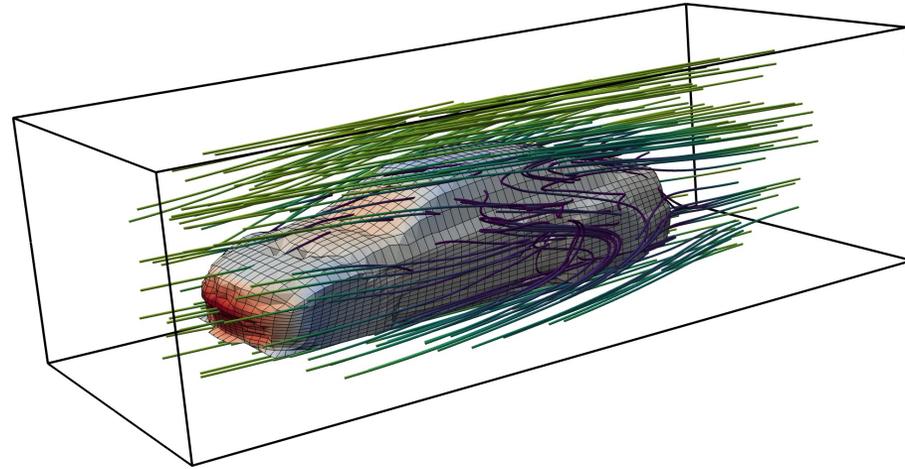
Car



Airplane



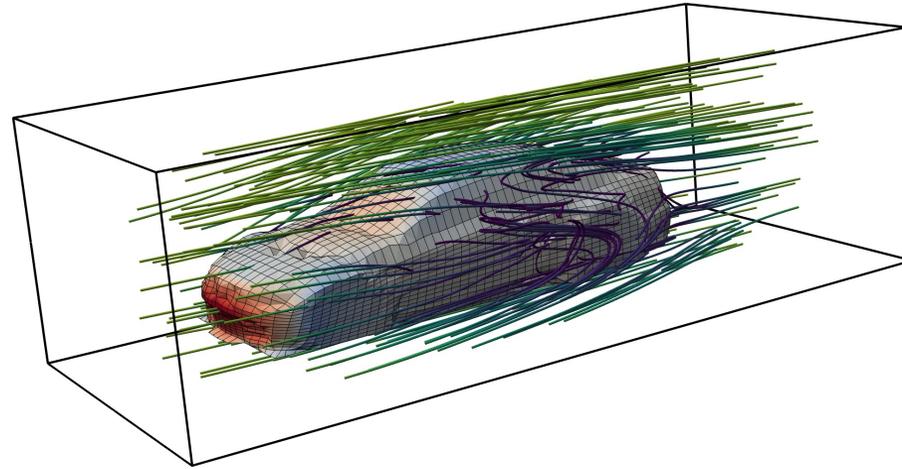
Challenges in Practical Industrial Design



Example: Estimate the drag coefficient of a given shape:

Surrounding Wind & Surface Pressure

Challenges in Practical Industrial Design

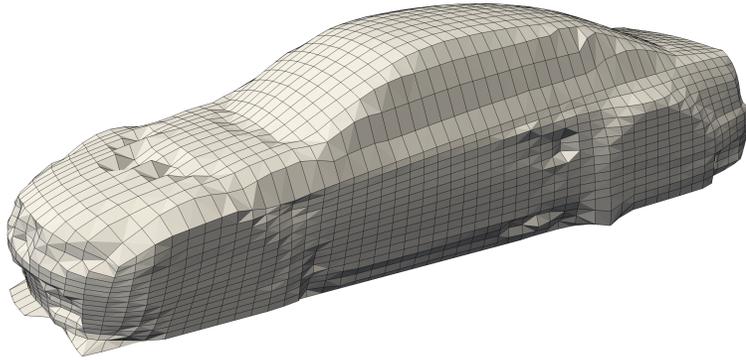


Example: Estimate the drag coefficient of a given shape:

Surrounding Wind & Surface Pressure

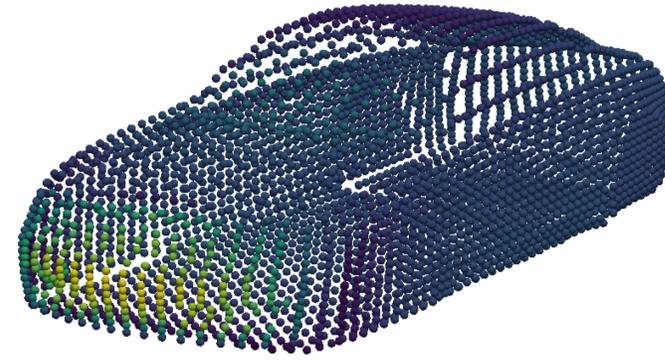
1. Large-scale meshes → **Huge computation cost**
2. Complex and unstructured geometrics → **Complex geometric learning**
3. Multiphysics interaction → **Intricate physical correlations**

Previous Work: Geometric Deep Learning



(1) Mesh

GraphSAGE, MeshGraphNet, etc

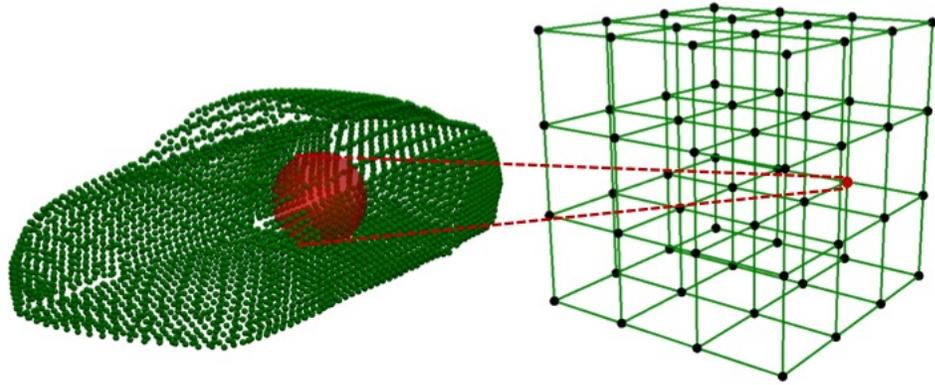


(2) Point Cloud

PointNet, Point Transformer, etc

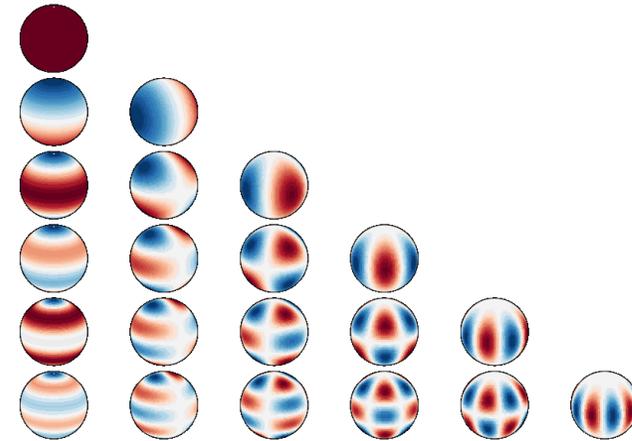
Excels in geometry modeling but fail in physics learning

Previous Work: Geometry-General Neural Operators



(1) GNN as Operators

GNO, GINO, etc

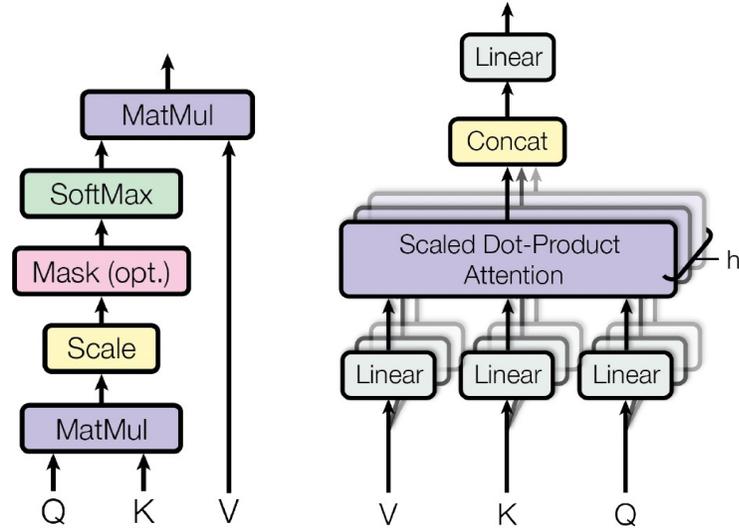
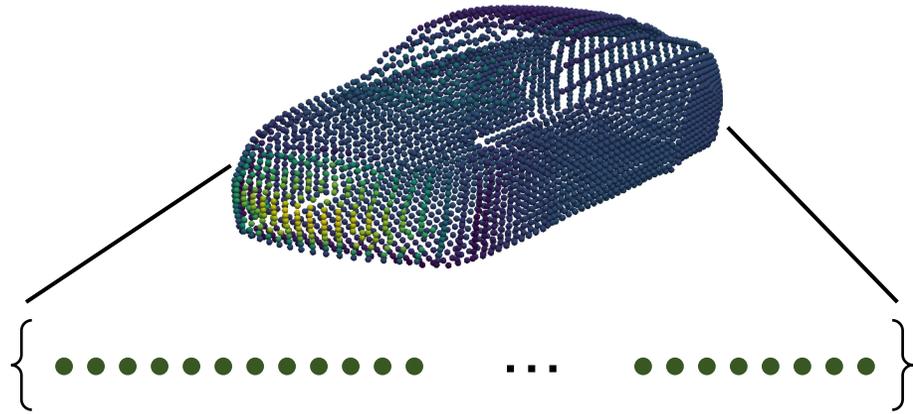


(2) FNO-Variants

geoFNO, SFNO, etc

Only focus on local physics or limited to periodic boundary

Transformer-based PDE Solvers

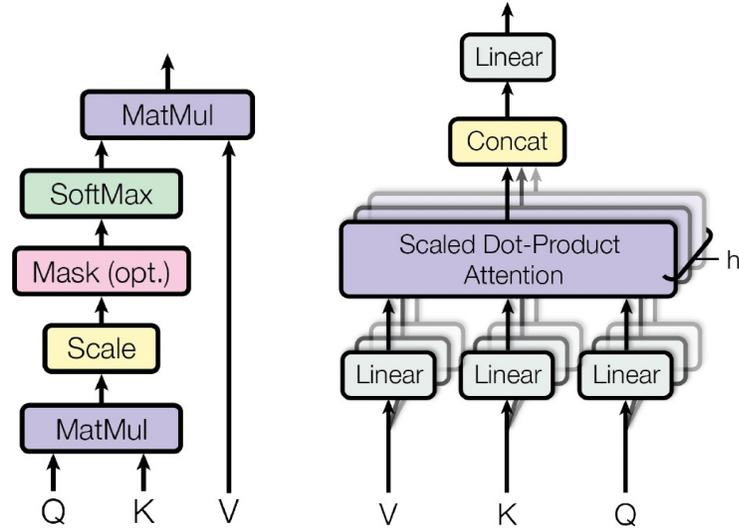
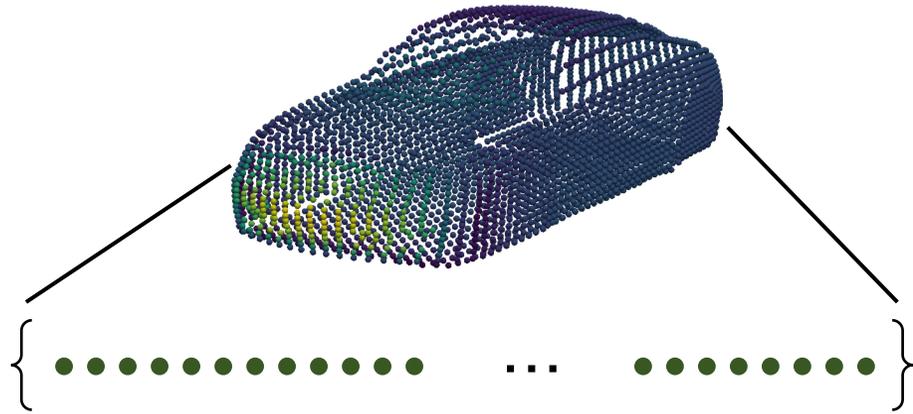


(1) Geometries as point sequences (2) Attention as Monte Carlo Integral

OFormer, Galerkin Transformer, etc

1. Quadratic complexity
2. Hard to capture physical correlations among massive points

Transformer-based PDE Solvers

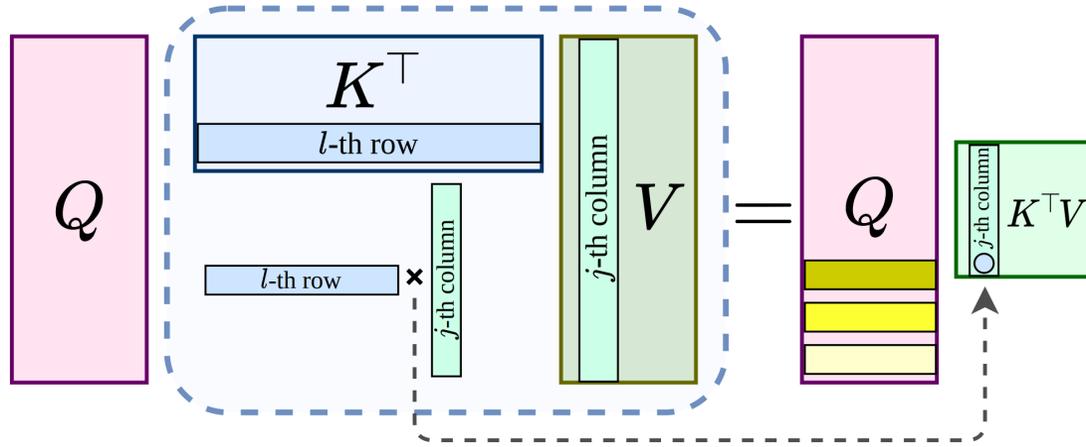


(1) Geometries as point sequences (2) Attention as Monte Carlo Integral

OFormer, Galerkin Transformer, etc

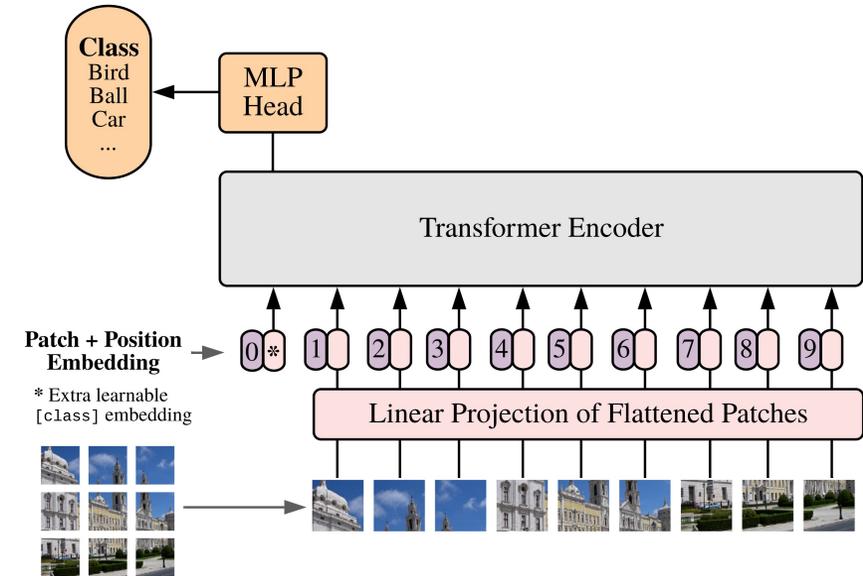
How to efficiently capture physical correlations underlying discretized meshes is the key to “transform” Transformers into practical PDE solvers

Related Work



(1) Linear Transformers

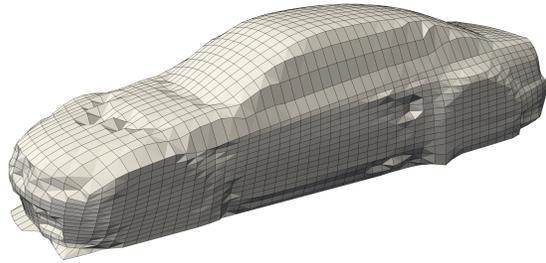
1. *Less informative attention*
2. *Individual points is insufficient for physics learning*



(2) Vision Transformer

- Augment features with patch ✓*
- Not applicable to irregular meshes*

A foundational Idea of Transolver



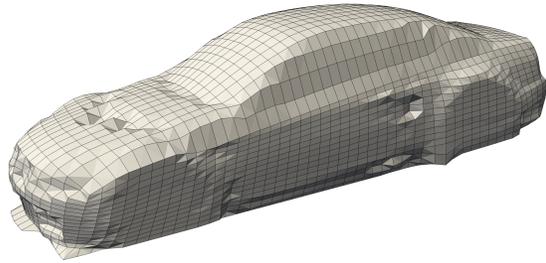
Discretized Domain

Previous Work

Being “trapped” to superficial and unwieldy meshes

Difficulties in Complexity, Geometry, Physics

A foundational Idea of Transolver

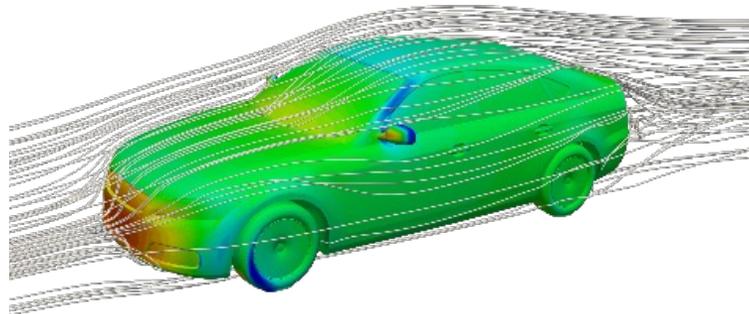


Discretized Domain

Previous Work

Being “trapped” to superficial and unwieldy meshes

Difficulties in Complexity, Geometry, Physics



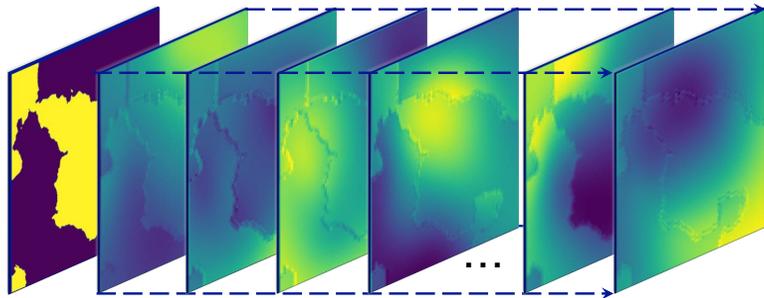
Physics Domain

Transolver

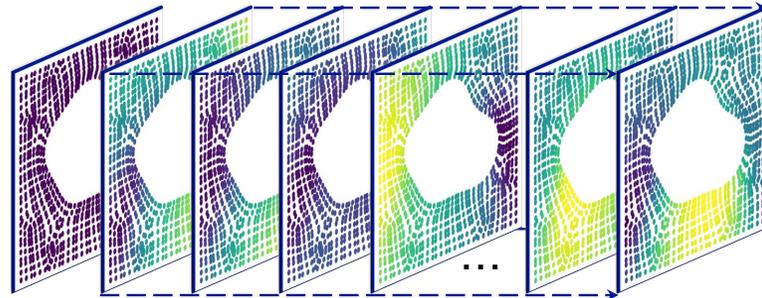
Learning **intrinsic physical states** under
complex and large-scale geometrics

Better Complexity, Geometry, Physics Modeling

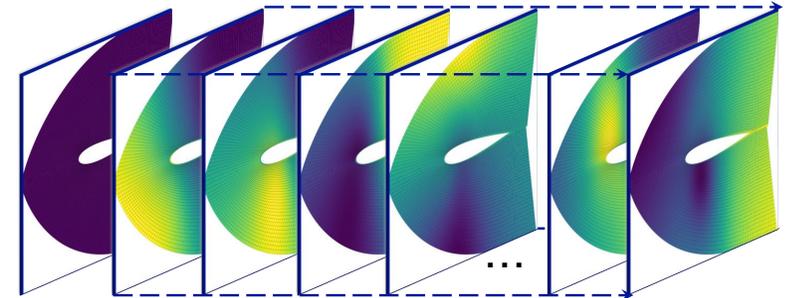
Learning Physical States



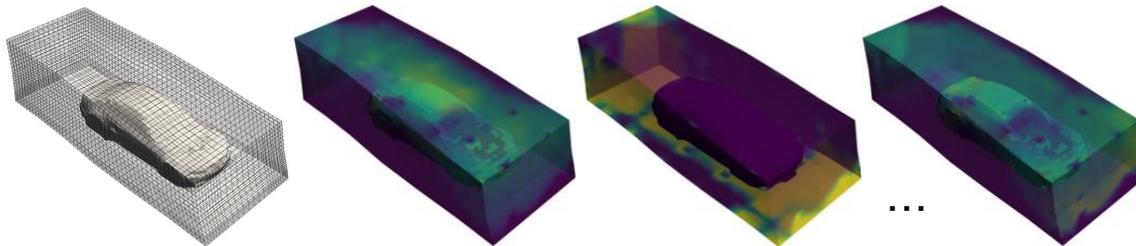
(a) Slices for Darcy, 2D Regular Grid



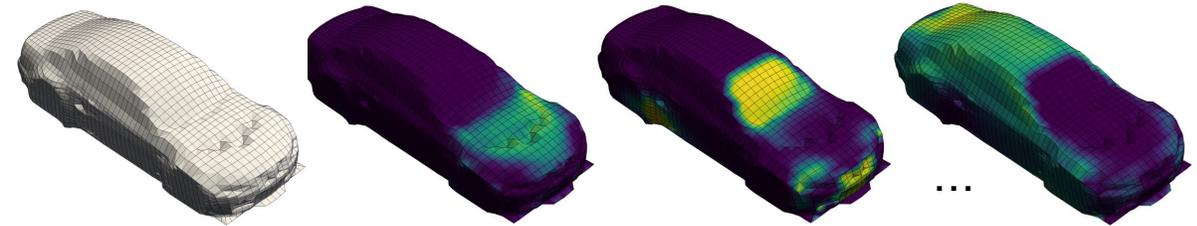
(b) Slices for Elasticity, 2D Point Cloud



(c) Slices for Airfoil, 2D Mesh



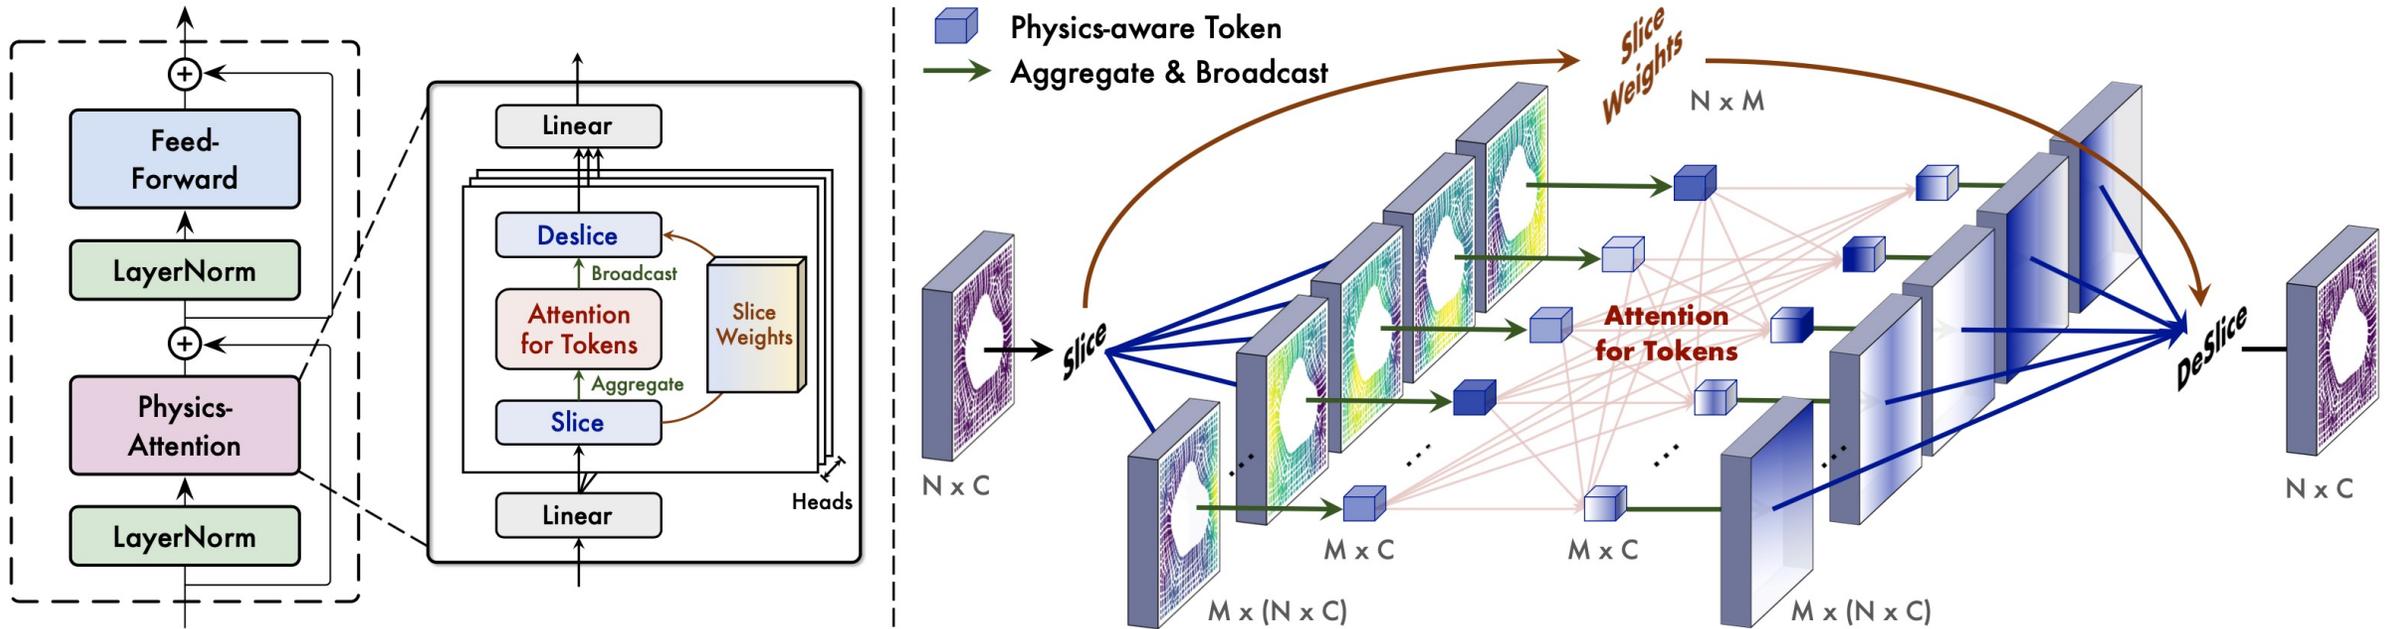
(d) Slices for Shape-Net Car Surrounding Velocity, 3D Volumes



(e) Slices for Shape-Net Car Surface Pressure, 3D Mesh

Mesh points under **similar physical states** will be ascribed to the same **slice** and then encoded into a physics-aware token.

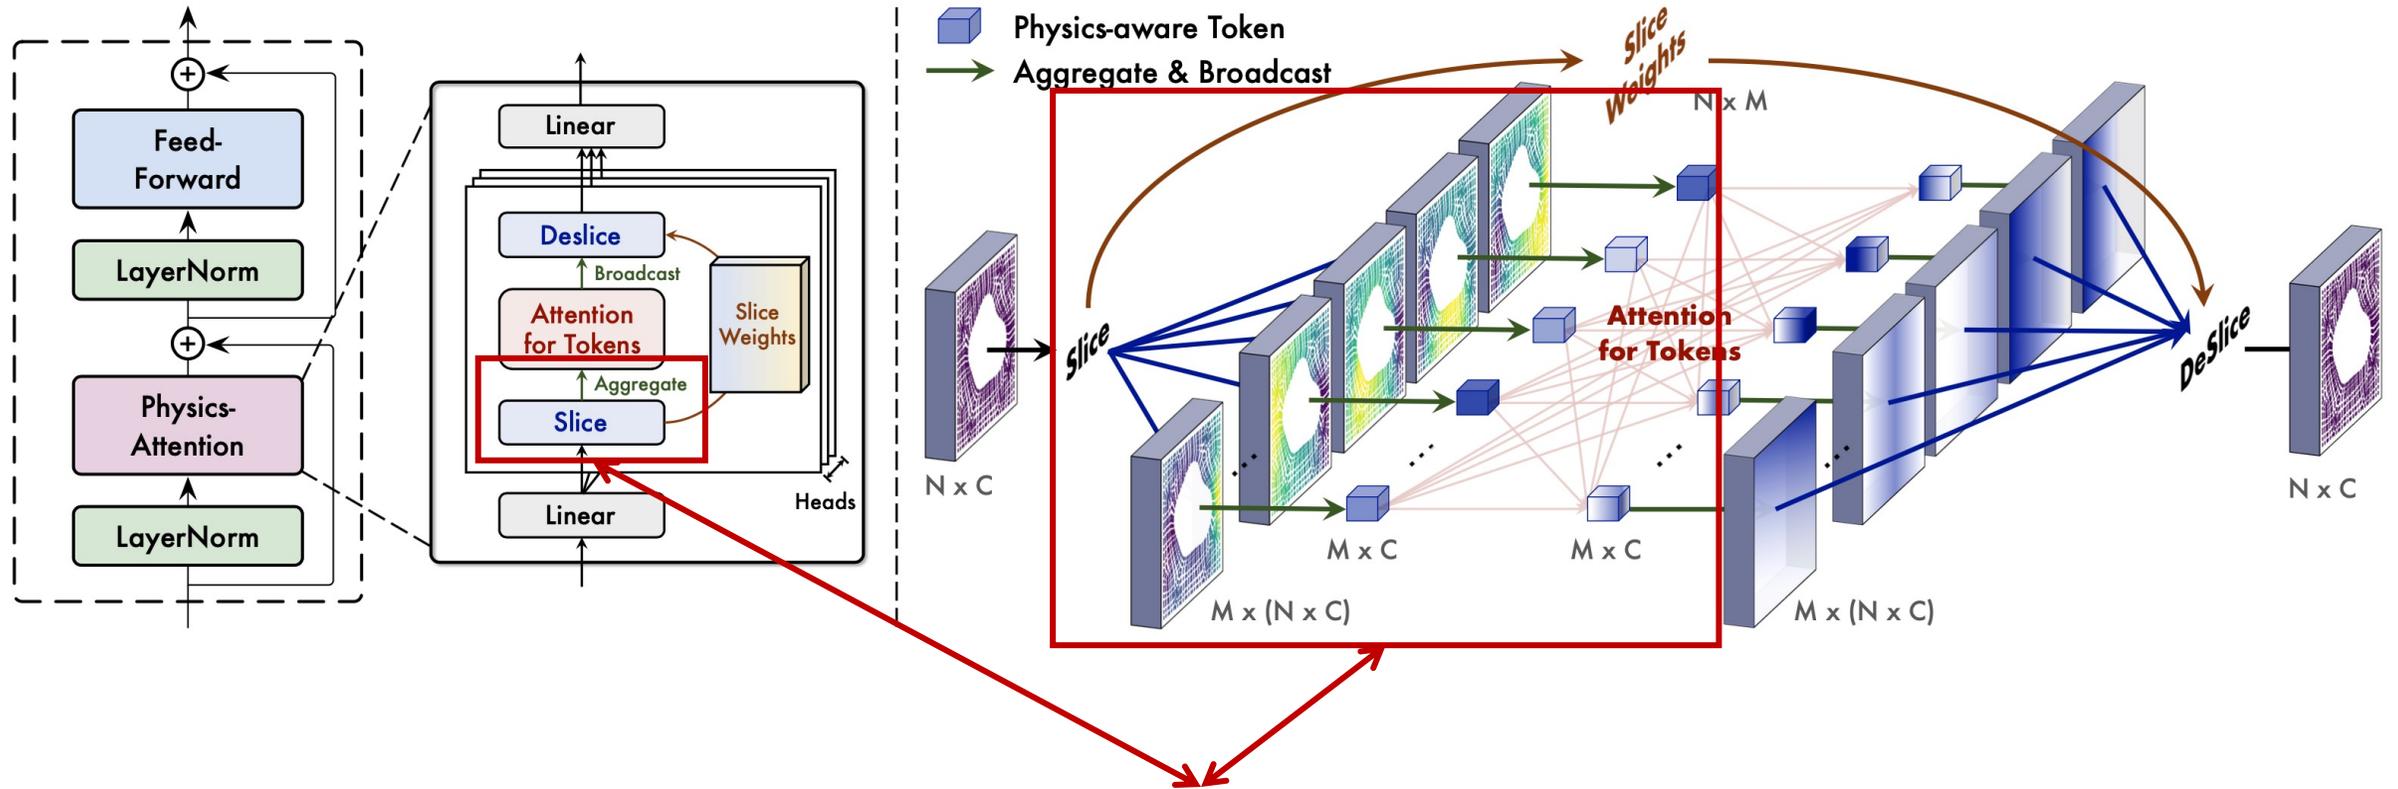
Overview of Transolver



Transolver applies attention to learned physical states (**Physics-Attention**)

- ① Mesh \rightarrow physics
- ② Attention (Integral)
- ③ Physics \rightarrow Mesh

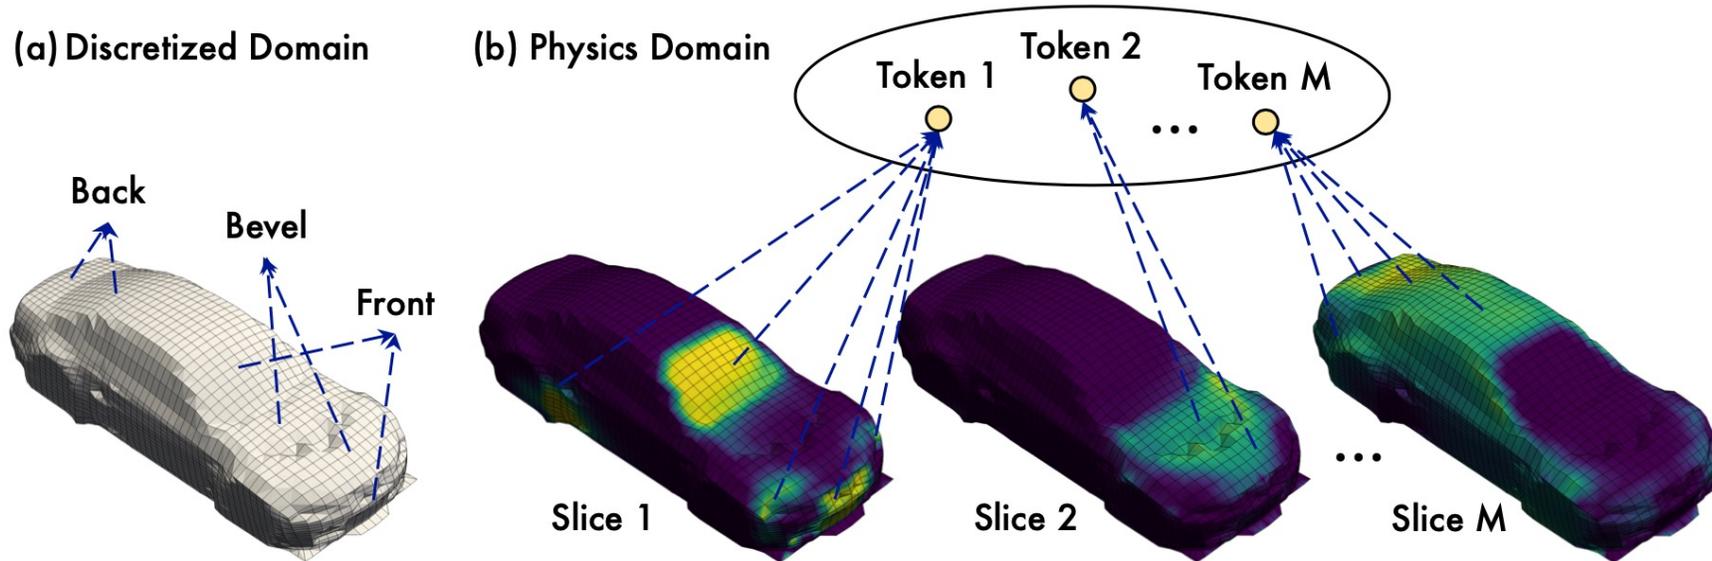
Overview of Transolver



① Mesh \rightarrow physics

To obtain physics-aware tokens

Mesh \rightarrow physics



1. Assign each point to slices with weights learned from features

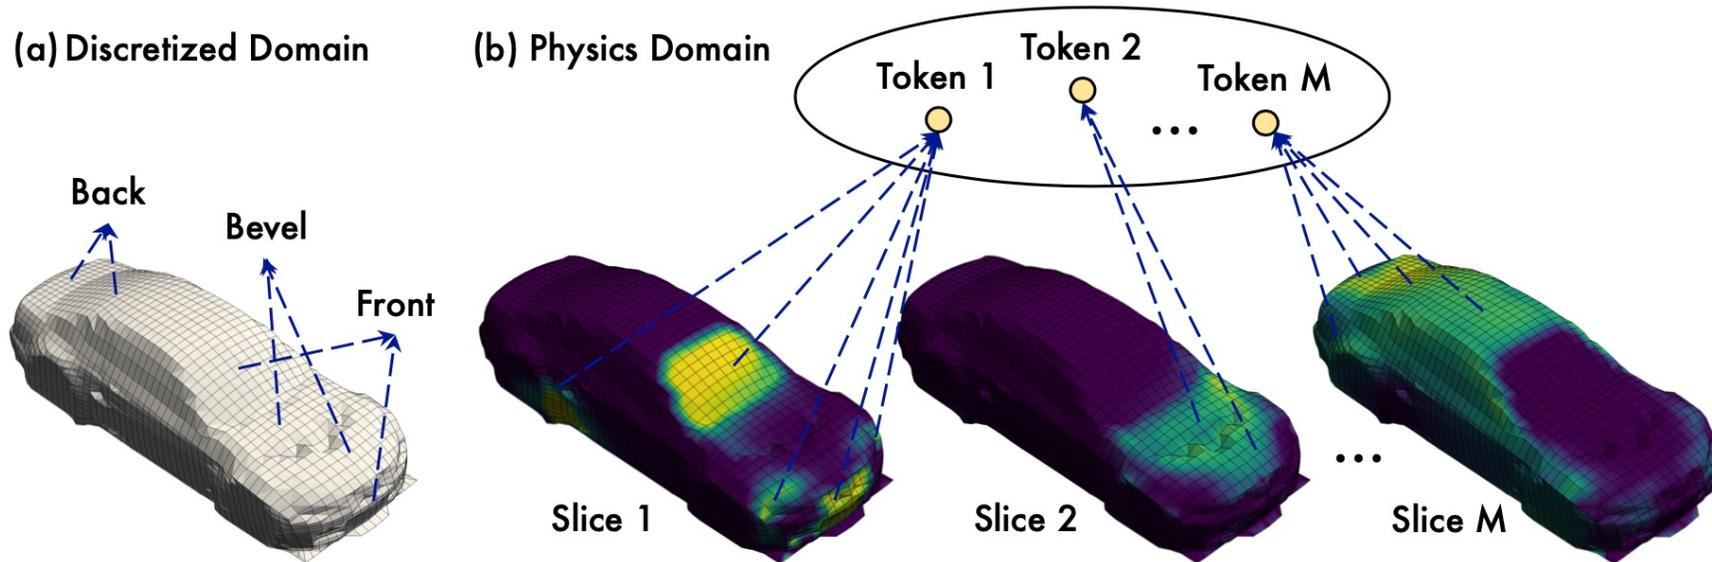
$$\{\mathbf{w}_i\}_{i=1}^N = \left\{ \underline{\text{Softmax}} \left(\text{Project}(\mathbf{x}_i) \right) \right\}_{i=1}^N$$

$$\mathbf{s}_j = \left\{ \mathbf{w}_{i,j} \mathbf{x}_i \right\}_{i=1}^N,$$

N Points to M Slices

Softmax for low-entropy slices

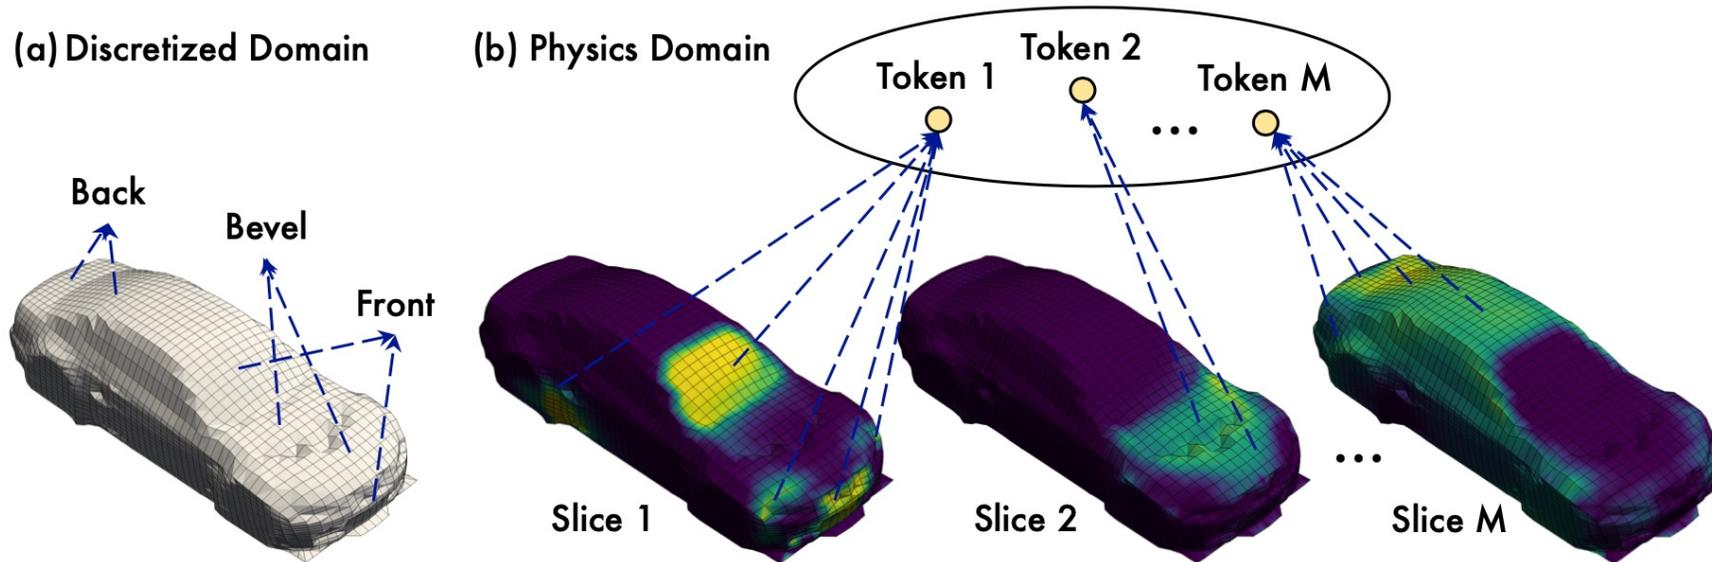
Mesh \rightarrow physics



1. Assign each point to slices
2. Aggregate slices for physics-aware tokens

$$\mathbf{z}_j = \frac{\sum_{i=1}^N \mathbf{s}_{j,i}}{\sum_{i=1}^N \mathbf{w}_{i,j}} = \frac{\sum_{i=1}^N \mathbf{w}_{i,j} \mathbf{x}_i}{\sum_{i=1}^N \mathbf{w}_{i,j}}$$

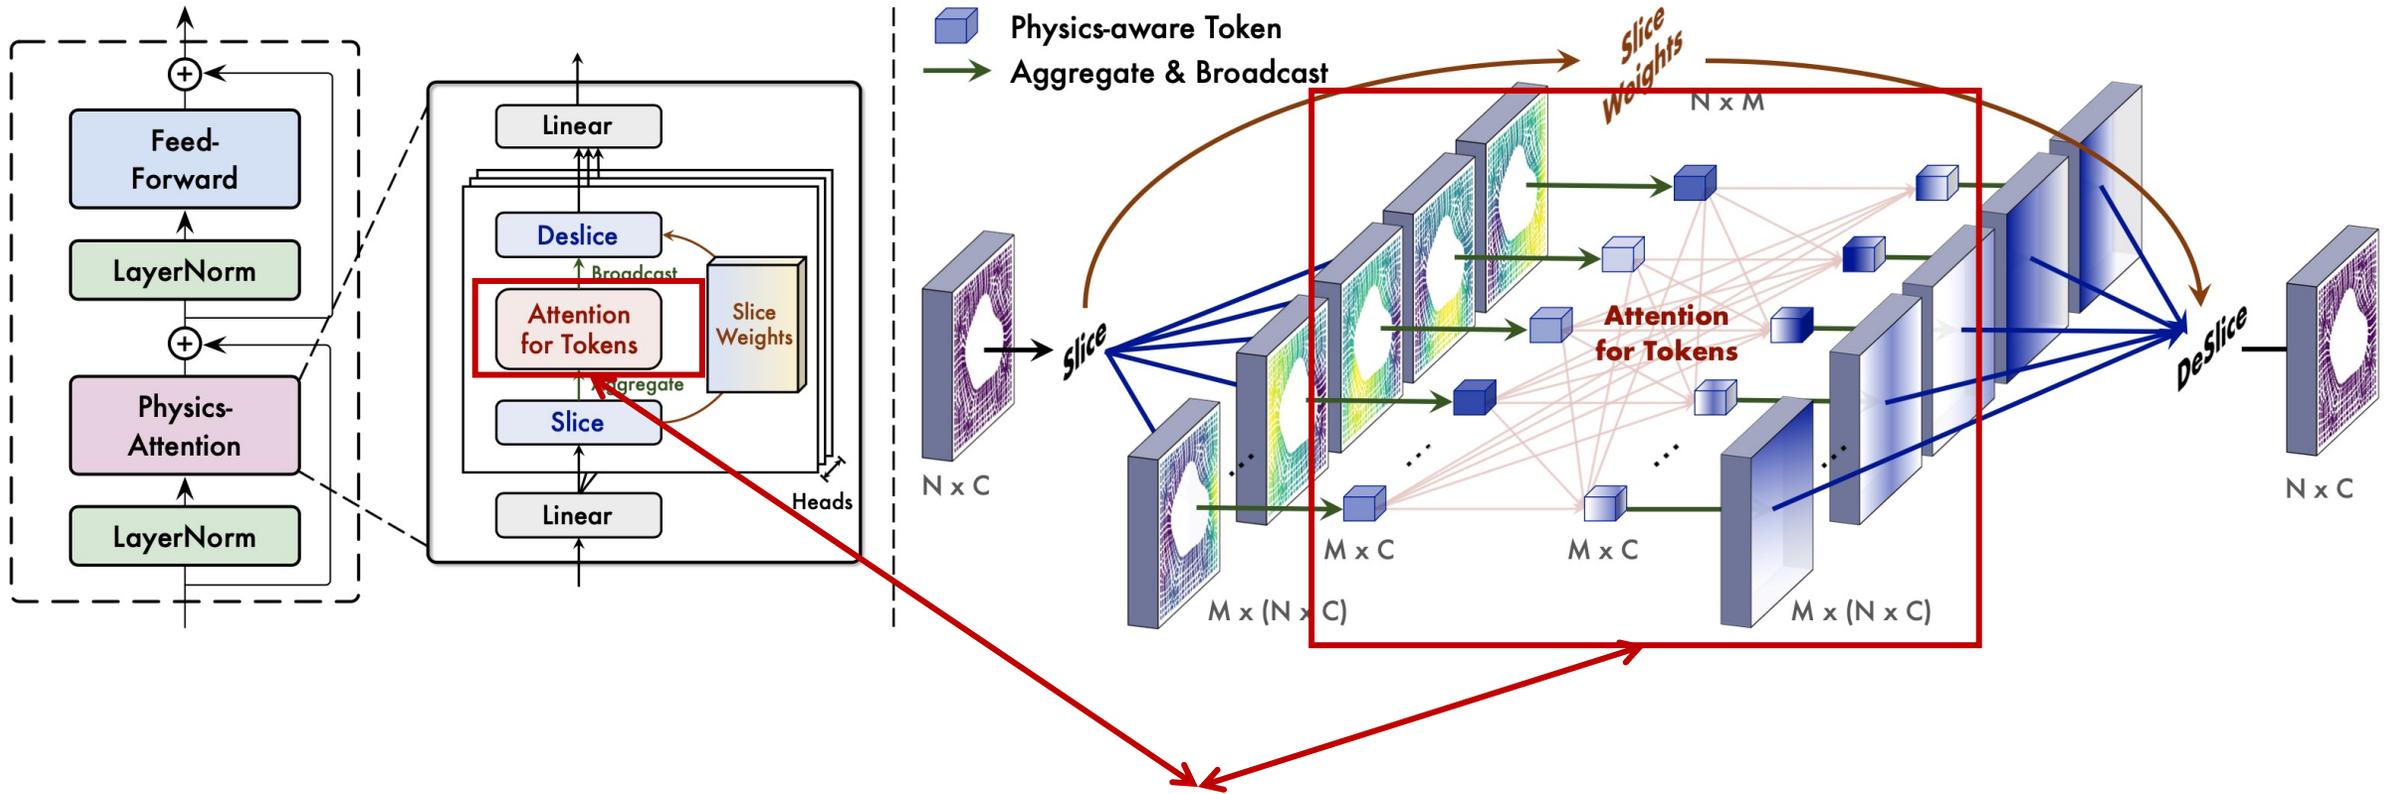
Mesh \rightarrow physics



1. Why slices can learn physically internal-consistent information
2. Learning slice is different from splitting computation area

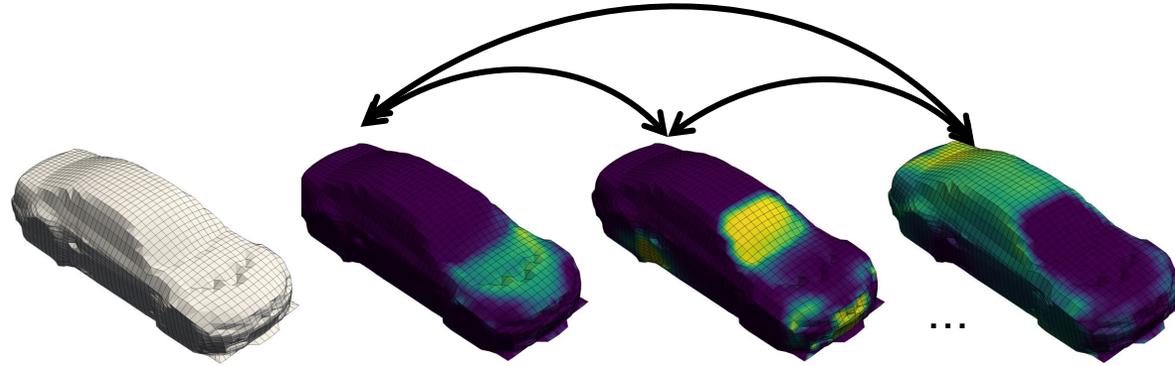
Ascribe physically similar but spatially distant points to the same slice

Overview of Transolver



② Attention among physics tokens
Approximate Integral to solve PDEs

Attention among physics tokens

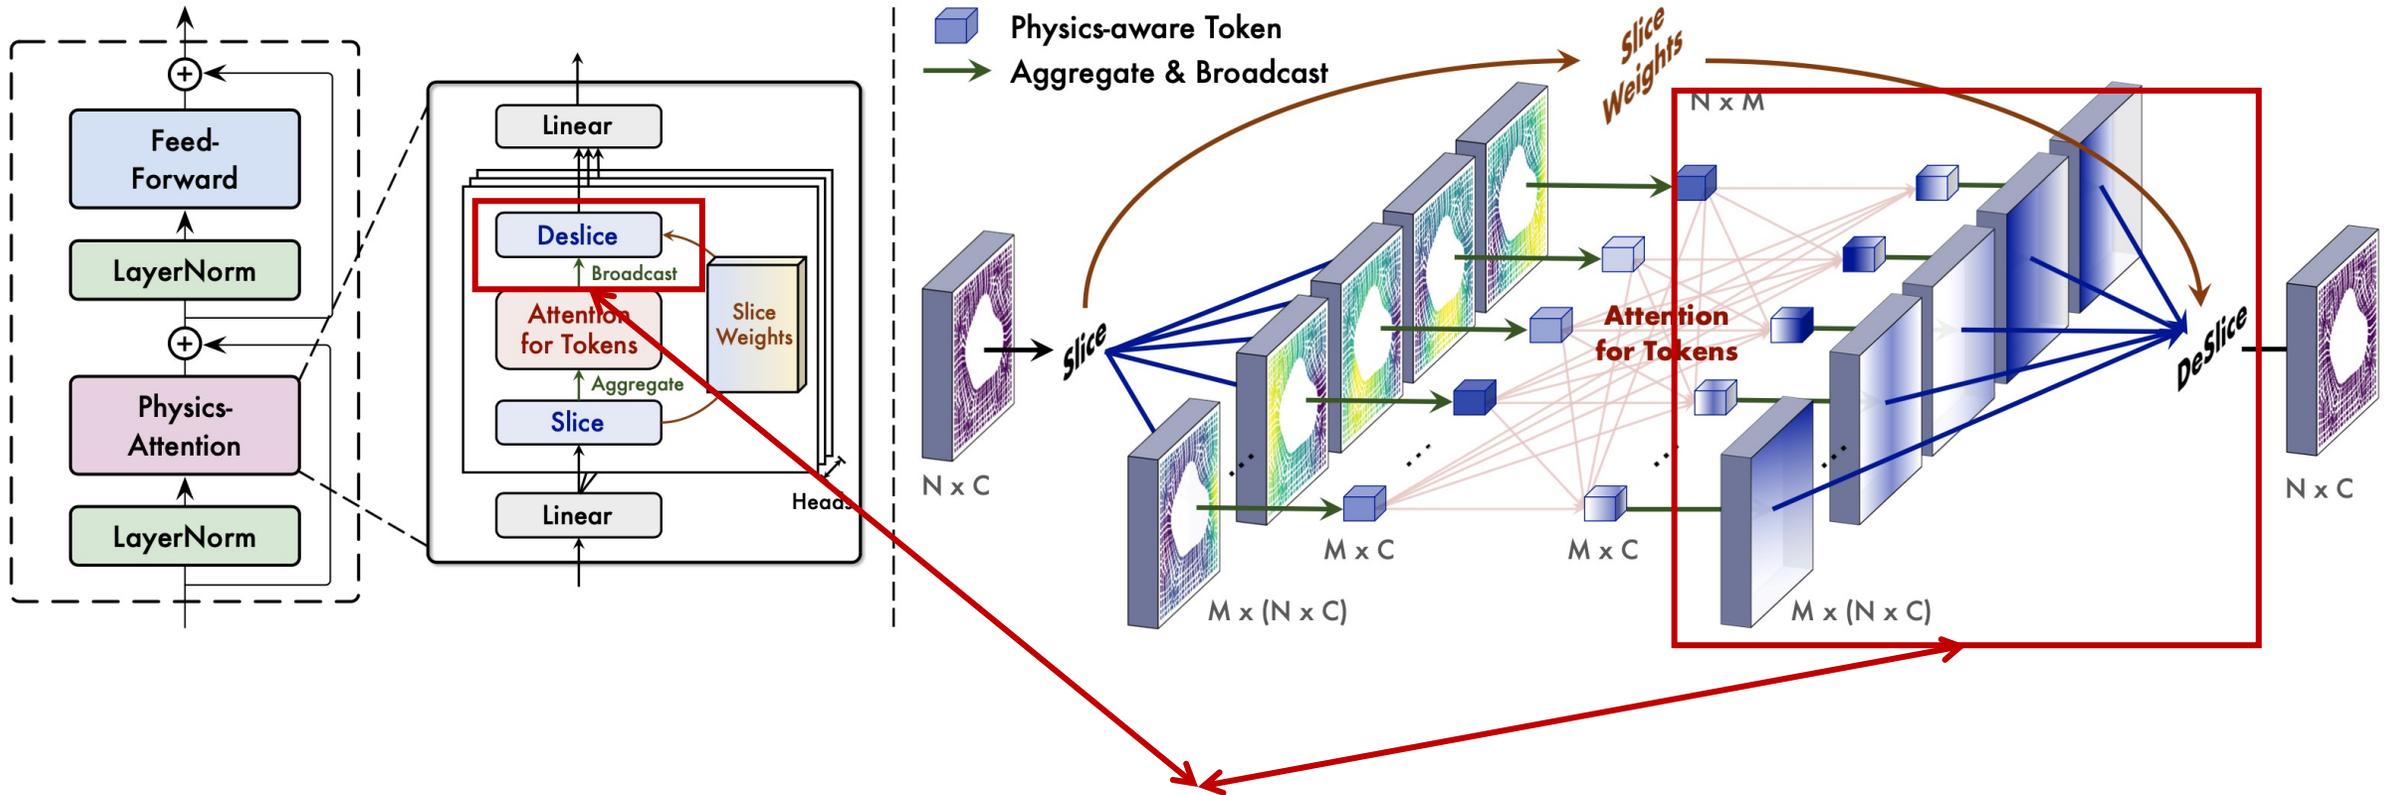


$$\mathbf{q}, \mathbf{k}, \mathbf{v} = \text{Linear}(\underline{\mathbf{z}}), \quad \mathbf{z}' = \text{Softmax} \left(\frac{\mathbf{q}\mathbf{k}^\top}{\sqrt{C}} \right) \mathbf{v}$$

Canonical attention among physics tokens

1. Complexity: $\mathcal{O}(N^2C) \rightarrow \mathcal{O}(M^2C)$
2. Capture interactions among physics states
3. Theorem: Attention as learnable integral operator

Overview of Transolver



③ Physics → Mesh

Project physics information back to mesh

$$\mathbf{x}'_i = \sum_{j=1}^M \mathbf{w}_{i,j} \mathbf{z}'_j$$

Theoretical Understanding of Transolver

1. Corollary of *Attention is a learnable integral*

Since attention mechanism is applied to tokens encoded from slices, **the step 2 (attention part of Transolver) is a learnable integral for the physics domain**

Is Physics-Attention still an input domain integral?

$$\mathcal{G}(\mathbf{u})(\mathbf{g}^*) = \int_{\Omega} \kappa(\mathbf{g}^*, \boldsymbol{\xi}) \mathbf{u}(\boldsymbol{\xi}) d\boldsymbol{\xi}$$

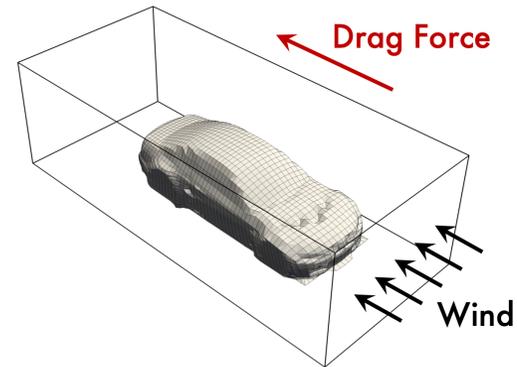
Theoretical Understanding of Transolver

$$\begin{aligned}
 \mathcal{G}(\mathbf{u})(\mathbf{g}) &= \int_{\Omega} \kappa(\mathbf{g}, \boldsymbol{\xi}) \mathbf{u}(\boldsymbol{\xi}) d\boldsymbol{\xi} \\
 &= \int_{\Omega_s} \kappa_{\text{ms}}(\mathbf{g}, \boldsymbol{\xi}_s) \mathbf{u}_s(\boldsymbol{\xi}_s) d\mathbf{g}^{-1}(\boldsymbol{\xi}_s) && (\kappa_{\text{ms}}(\cdot, \cdot) : \Omega \times \Omega_s \rightarrow \mathbb{R}^{C \times C} \text{ is a kernel function}) \\
 &= \int_{\Omega_s} \kappa_{\text{ms}}(\mathbf{g}, \boldsymbol{\xi}_s) \mathbf{u}_s(\boldsymbol{\xi}_s) |\det(\nabla_{\boldsymbol{\xi}_s} \mathbf{g}^{-1}(\boldsymbol{\xi}_s))| d\boldsymbol{\xi}_s \\
 &= \int_{\Omega_s} \left(\frac{\int_{\Omega_s} w_{\mathbf{g}, \boldsymbol{\xi}'_s} \kappa_{\text{ss}}(\boldsymbol{\xi}'_s, \boldsymbol{\xi}_s) d\boldsymbol{\xi}'_s}{\int_{\Omega_s} w_{\mathbf{g}, \boldsymbol{\xi}'_s} d\boldsymbol{\xi}'_s} \right) \mathbf{u}_s(\boldsymbol{\xi}_s) |\det(\nabla_{\boldsymbol{\xi}_s} \mathbf{g}^{-1}(\boldsymbol{\xi}_s))| d\boldsymbol{\xi}_s && (\kappa_{\text{ms}} \text{ is a linear combination of } \kappa_{\text{ss}} \text{ with weights } w_{*,*}) \\
 &= \int_{\Omega_s} \underbrace{w_{\mathbf{g}, \boldsymbol{\xi}'_s}}_{\text{DeSlice}} \int_{\Omega_s} \underbrace{\kappa_{\text{ss}}(\boldsymbol{\xi}'_s, \boldsymbol{\xi}_s)}_{\text{Attention among slice tokens}} \underbrace{\mathbf{u}_s(\boldsymbol{\xi}_s)}_{\text{Slice token}} |\det(\nabla_{\boldsymbol{\xi}_s} \mathbf{g}^{-1}(\boldsymbol{\xi}_s))| d\boldsymbol{\xi}_s d\boldsymbol{\xi}'_s && (\text{Suppose that } \int_{\Omega_s} w_{\mathbf{g}, \boldsymbol{\xi}'_s} d\boldsymbol{\xi}'_s = 1) \\
 &\approx \underbrace{\sum_{j=1}^M \mathbf{w}_{i,j}}_{\text{Eq. (4)}} \underbrace{\sum_{t=1}^M \frac{\exp\left(\left(\mathbf{W}_{\mathbf{q}} \mathbf{u}_s(\boldsymbol{\xi}_{s,j})\right) \left(\mathbf{W}_{\mathbf{k}} \mathbf{u}_s(\boldsymbol{\xi}_{s,t})\right)^{\top} / \tau\right)}{\sum_{p=1}^M \exp\left(\left(\mathbf{W}_{\mathbf{q}} \mathbf{u}_s(\boldsymbol{\xi}_{s,j})\right) \left(\mathbf{W}_{\mathbf{k}} \mathbf{u}_s(\boldsymbol{\xi}_{s,p})\right)^{\top} / \tau\right)}}_{\text{Eq. (3)}} \underbrace{\mathbf{W}_{\mathbf{v}} \left(\frac{\sum_{p=1}^N \mathbf{w}_{p,t} \mathbf{u}(\mathbf{g}_p)}{\sum_{p=1}^N \mathbf{w}_{p,t}} \right)}_{\text{Eq. (2)}} && (\text{Lemma A.1}) \\
 &= \sum_{j=1}^M \mathbf{w}_{i,j} \sum_{t=1}^M \frac{\exp(\mathbf{q}_j \mathbf{k}_t^{\top} / \tau)}{\sum_{p=1}^M \exp(\mathbf{q}_j \mathbf{k}_p^{\top} / \tau)} \mathbf{v}_t,
 \end{aligned}$$

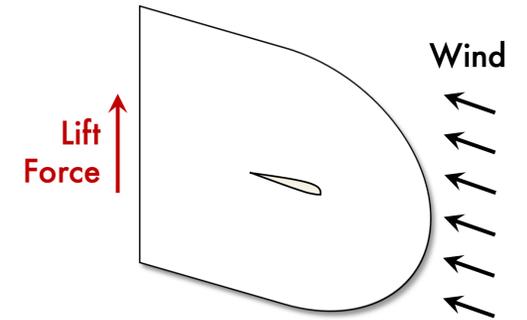
All the designs in Transolver can be directly derived.

Experiments

GEOMETRY	BENCHMARKS	#DIM	#MESH
POINT CLOUD	ELASTICITY	2D	972
STRUCTURED MESH	PLASTICITY	2D+TIME	3,131
	AIRFOIL	2D	11,271
	PIPE	2D	16,641
REGULAR GRID	NAVIER-STOKES	2D+TIME	4,096
	DARCY	2D	7,225
UNSTRUCTURED MESH	SHAPE-NET CAR	3D	32,186
	AIRFRANS	2D	32,000



(a) Shape-Net Car



(b) AirFRANS

Six standard benchmarks, two practical design tasks

More than 20 baselines

Standard PDE-Solving Benchmarks

MODEL	POINT CLOUD	STRUCTURED MESH			REGULAR GRID	
	ELASTICITY	PLASTICITY	AIRFOIL	PIPE	NAVIER-STOKES	DARCY
FNO (LI ET AL., 2021)	/	/	/	/	0.1556	0.0108
WMT (GUPTA ET AL., 2021)	0.0359	0.0076	0.0075	0.0077	0.1541	0.0082
U-FNO (WEN ET AL., 2022)	0.0239	0.0039	0.0269	0.0056	0.2231	0.0183
GEO-FNO (LI ET AL., 2022)	0.0229	0.0074	0.0138	0.0067	0.1556	0.0108
U-NO (RAHMAN ET AL., 2023)	0.0258	0.0034	0.0078	0.0100	0.1713	0.0113
F-FNO (TRAN ET AL., 2023)	0.0263	0.0047	0.0078	0.0070	0.2322	0.0077
LSM (WU ET AL., 2023)	0.0218	0.0025	<u>0.0059</u>	0.0050	0.1535	<u>0.0065</u>
GALERKIN (CAO, 2021)	0.0240	0.0120	0.0118	0.0098	0.1401	0.0084
HT-NET (LIU ET AL., 2022)	/	0.0333	0.0065	0.0059	0.1847	0.0079
OFORMER (LI ET AL., 2023C)	0.0183	<u>0.0017</u>	0.0183	0.0168	0.1705	0.0124
GNOT (HAO ET AL., 2023)	<u>0.0086</u>	<u>0.0336</u>	0.0076	<u>0.0047</u>	0.1380	0.0105
FACTFORMER (LI ET AL., 2023D)	/	0.0312	0.0071	0.0060	0.1214	0.0109
ONO (XIAO ET AL., 2024)	0.0118	0.0048	0.0061	0.0052	<u>0.1195</u>	0.0076
TRANSOLVER (OURS)	0.0064	0.0012	0.0053	0.0033	0.0900	0.0057
RELATIVE PROMOTION	25.6%	29.4%	10.2%	29.7%	24.7%	12.3%

Transolver achieves 22% error reduction over the second-best model

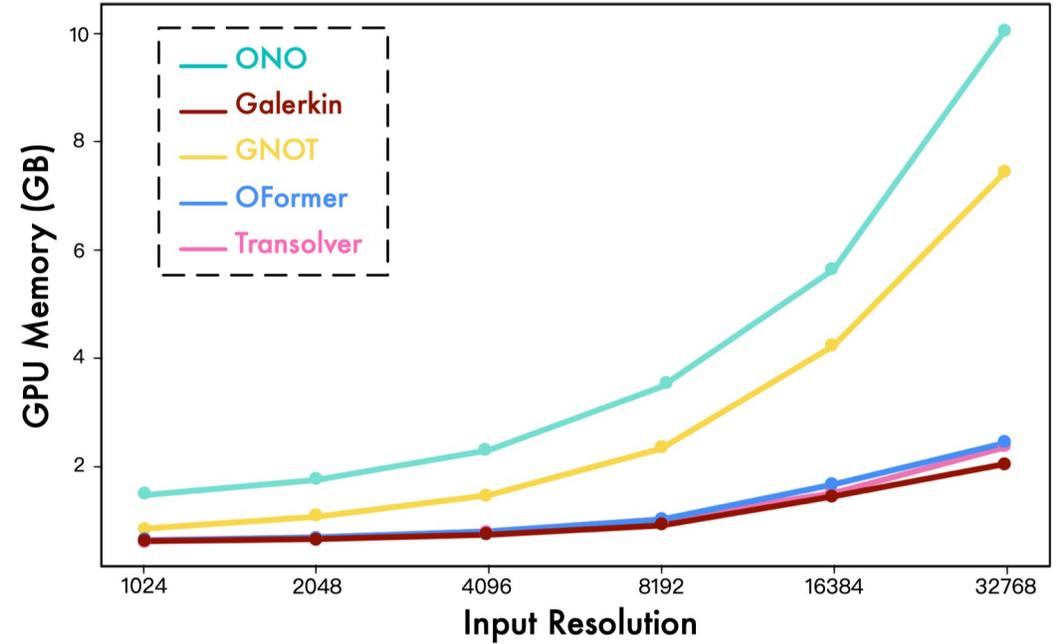
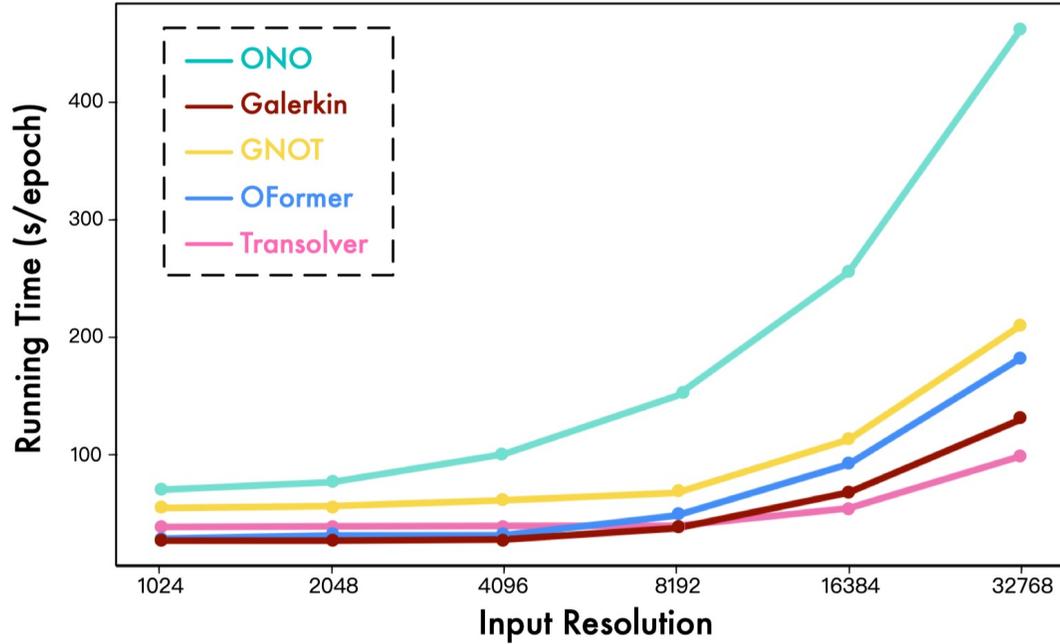
Practical Design Tasks

MODEL*	SHAPE-NET CAR				AIRFRANS			
	VOLUME ↓	SURF ↓	C_D ↓	ρ_D ↑	VOLUME ↓	SURF ↓	C_L ↓	ρ_L ↑
SIMPLE MLP	0.0512	0.1304	0.0307	0.9496	0.0081	0.0200	0.2108	0.9932
GRAPHSAGE (HAMILTON ET AL., 2017)	0.0461	0.1050	0.0270	0.9695	0.0087	0.0184	0.1476	0.9964
POINTNET (QI ET AL., 2017)	0.0494	0.1104	0.0298	0.9583	0.0253	0.0996	0.1973	0.9919
GRAPH U-NET (GAO & JI, 2019)	0.0471	0.1102	0.0226	0.9725	0.0076	0.0144	0.1677	0.9949
MESHGRAPHNET (PFAFF ET AL., 2021)	0.0354	0.0781	0.0168	0.9840	0.0214	0.0387	0.2252	0.9945
GNO (LI ET AL., 2020A)	0.0383	0.0815	0.0172	0.9834	0.0269	0.0405	0.2016	0.9938
GALERKIN (CAO, 2021)	0.0339	0.0878	0.0179	0.9764	0.0074	0.0159	0.2336	0.9951
GEO-FNO (LI ET AL., 2022)	0.1670	0.2378	0.0664	0.8280	0.0361	0.0301	0.6161	0.9257
GNOT (HAO ET AL., 2023)	0.0329	0.0798	0.0178	0.9833	0.0049	0.0152	0.1992	0.9942
GINO (LI ET AL., 2023A)	0.0386	0.0810	0.0184	0.9826	0.0297	0.0482	0.1821	0.9958
3D-GEOCA (DENG ET AL., 2024)	0.0319	0.0779	0.0159	0.9842	/	/	/	/
TRANSOLVER (OURS)	0.0207	0.0745	0.0103	0.9935	0.0037	0.0142	0.1030	0.9978

Design-oriented metrics: Drag/lift coefficients and their Spearman's correlation

Transolver performs best in both physics and design-oriented metrics

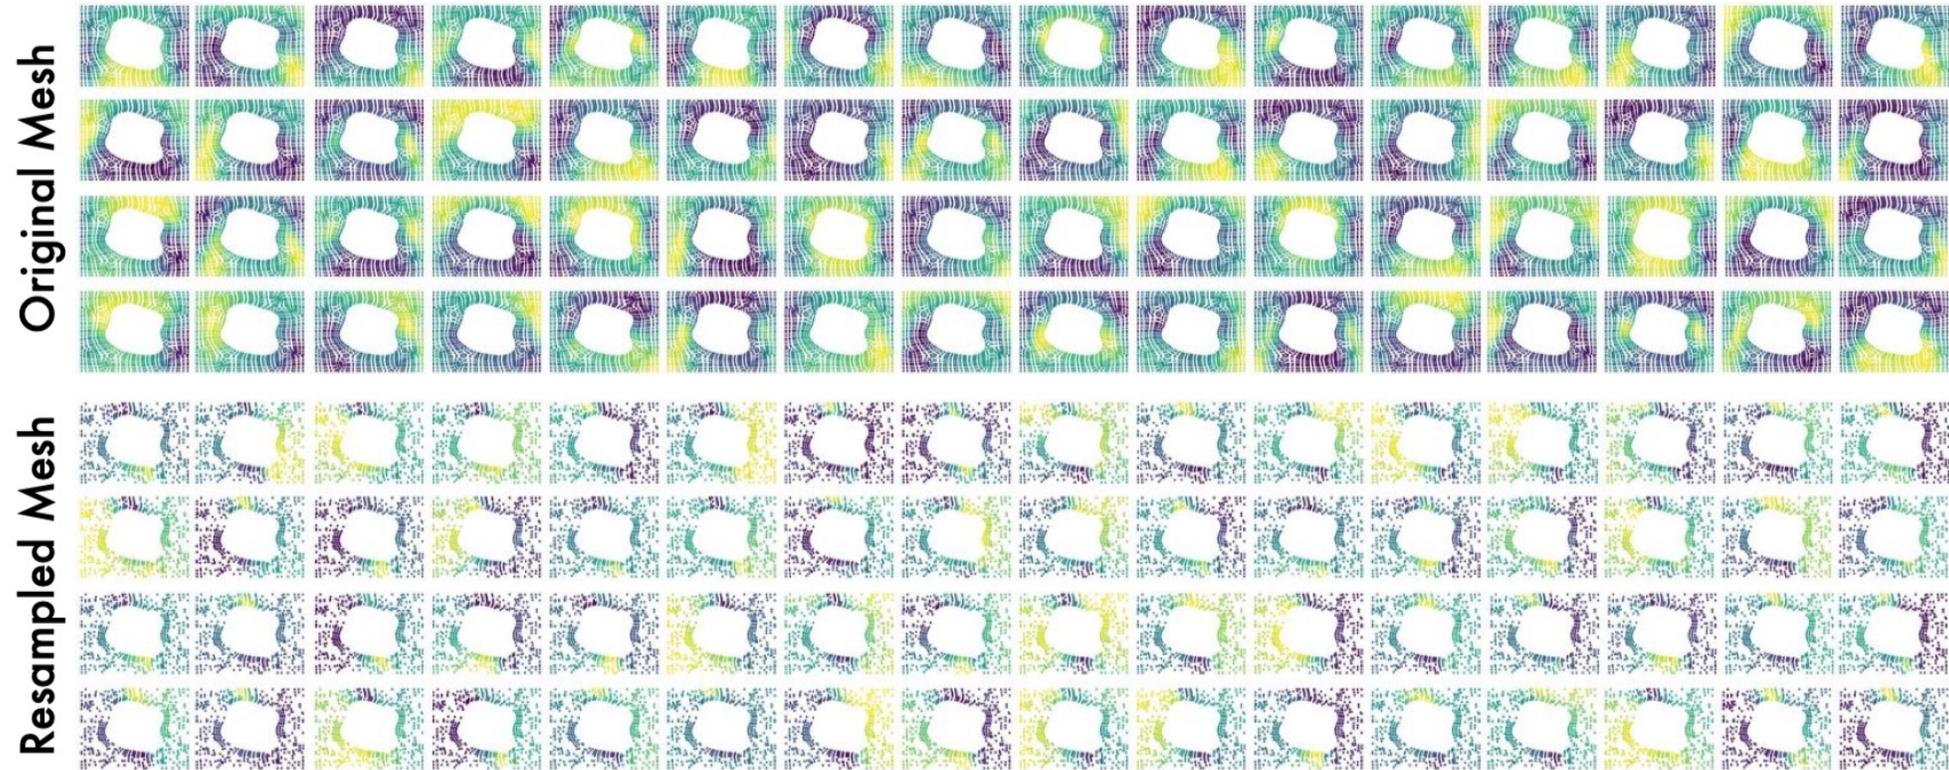
Efficiency



Favorable efficiency and performance balance

Transolver is faster than linear Transformers in large-scale meshes.

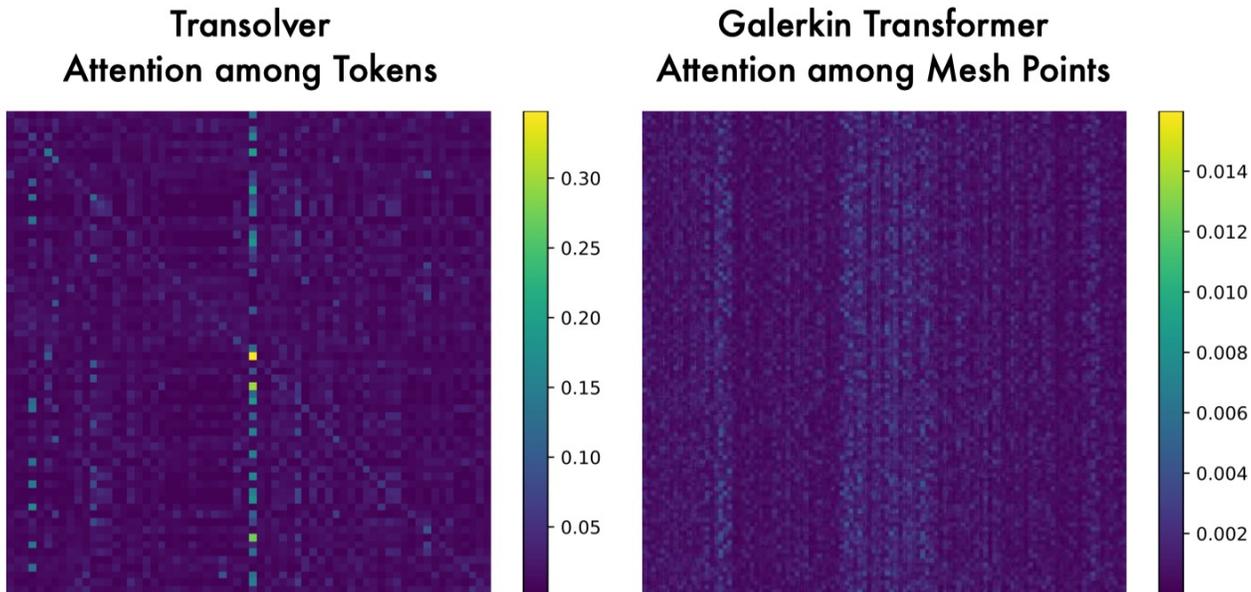
Physics-Attention Visualization



Slice visualization on Elasticity

Transolver is mesh-free, precisely captures states **even on broken meshes**

Physics-Attention Visualization

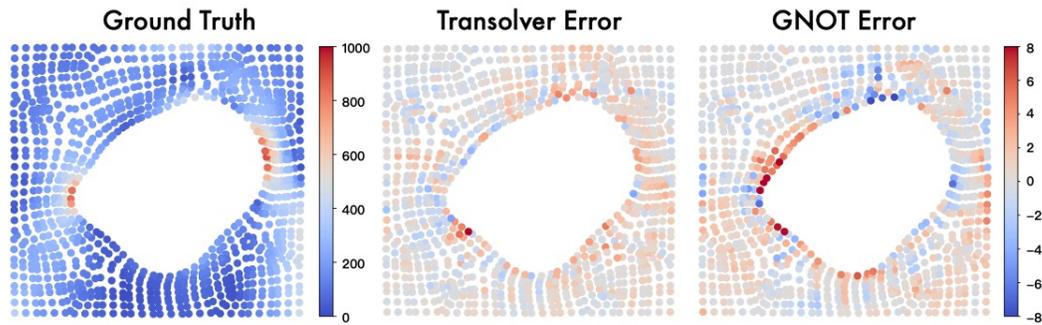


Kullback–Leibler (KL) divergence between attention weights and uniform distribution

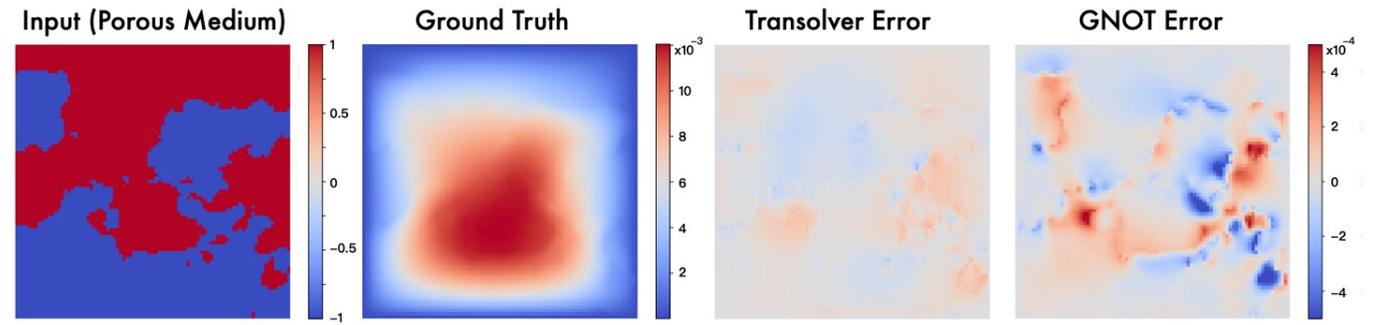
BENCHMARKS	GALERKIN (CAO, 2021)	TRANSOLVER (OURS)
ELASTICITY (972 MESH POINTS)	0.3803	1.7795
DARCY (7,225 MESH POINTS)	0.2739	1.8274

Physics-Attention can learn **more informative physical correlations**

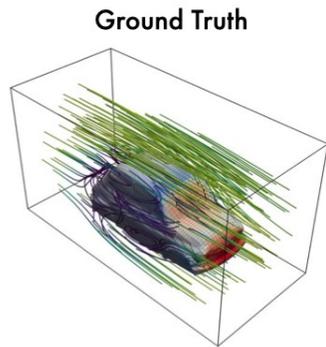
Showcases



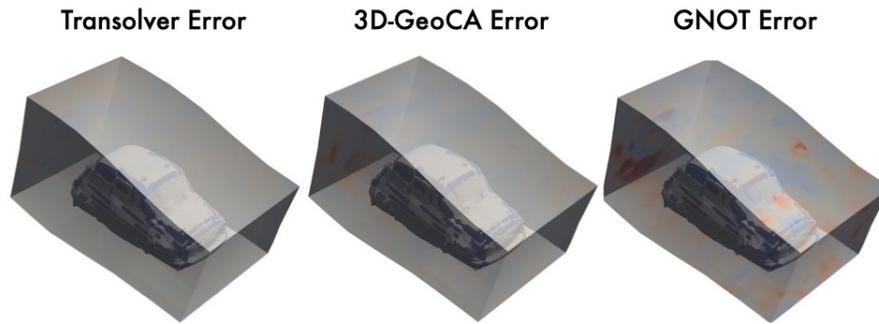
(a) Elasticity



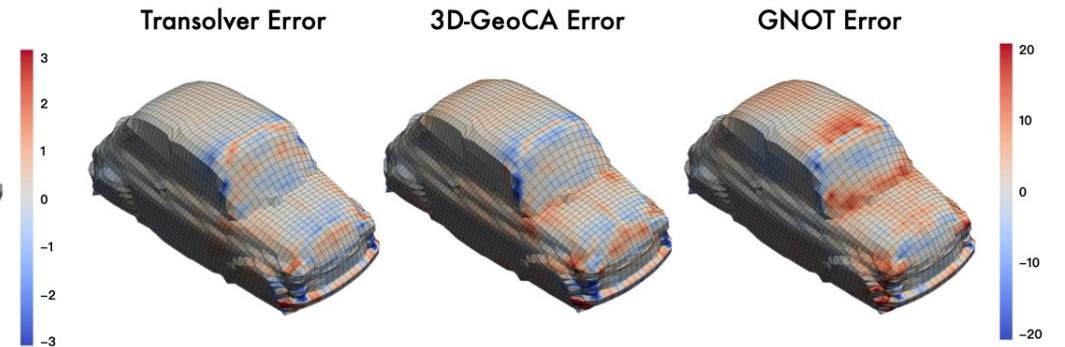
(b) Darcy



(c) Shape-Net Car



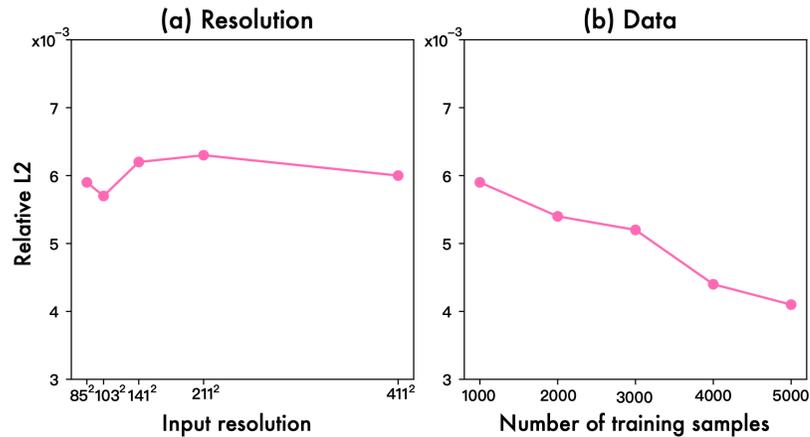
(c.1) Surrounding Velocity



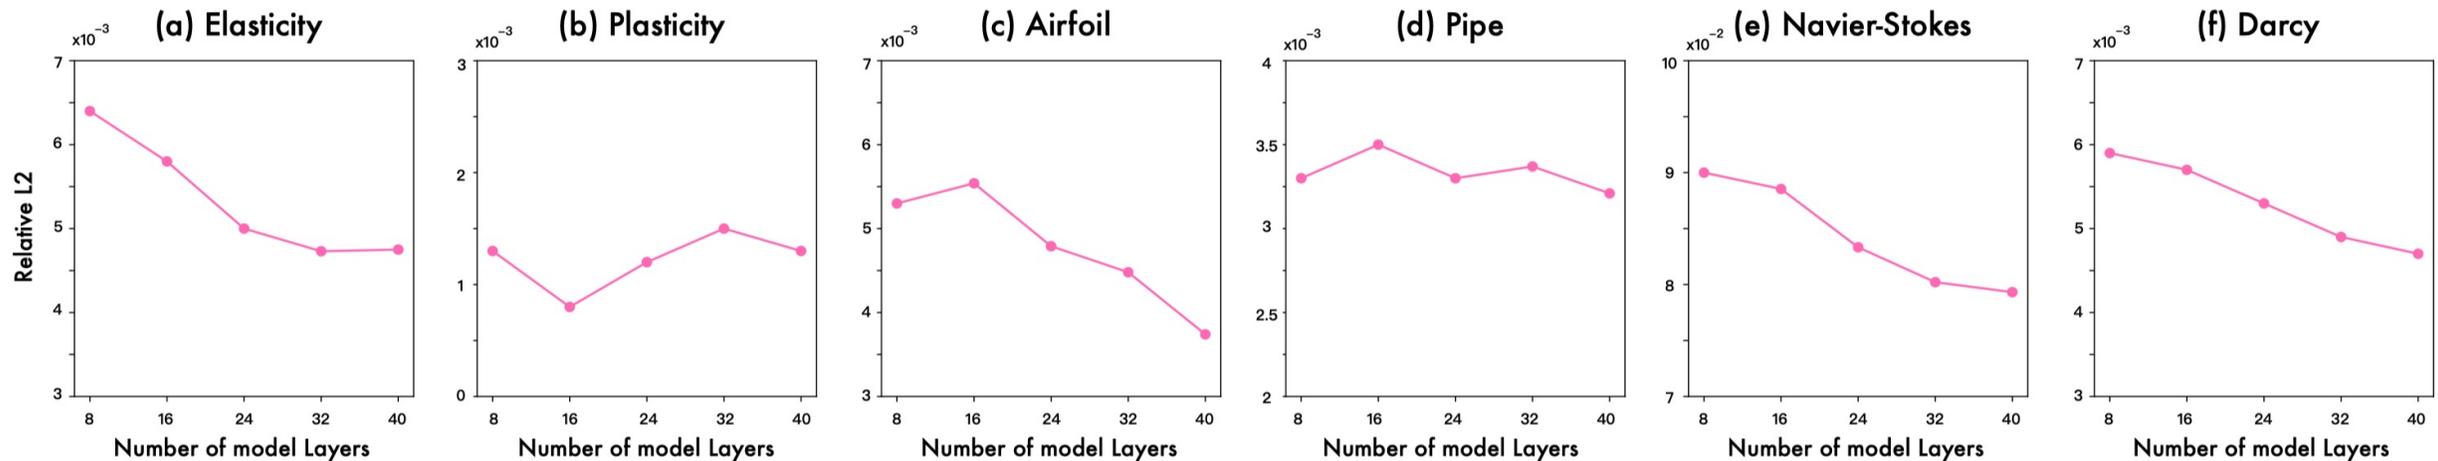
(c.2) Surface Pressure

Transolver excels in solving **multiphysics PDEs on hybrid geometrics**

Pursuing PDE Foundation Models: Scalability

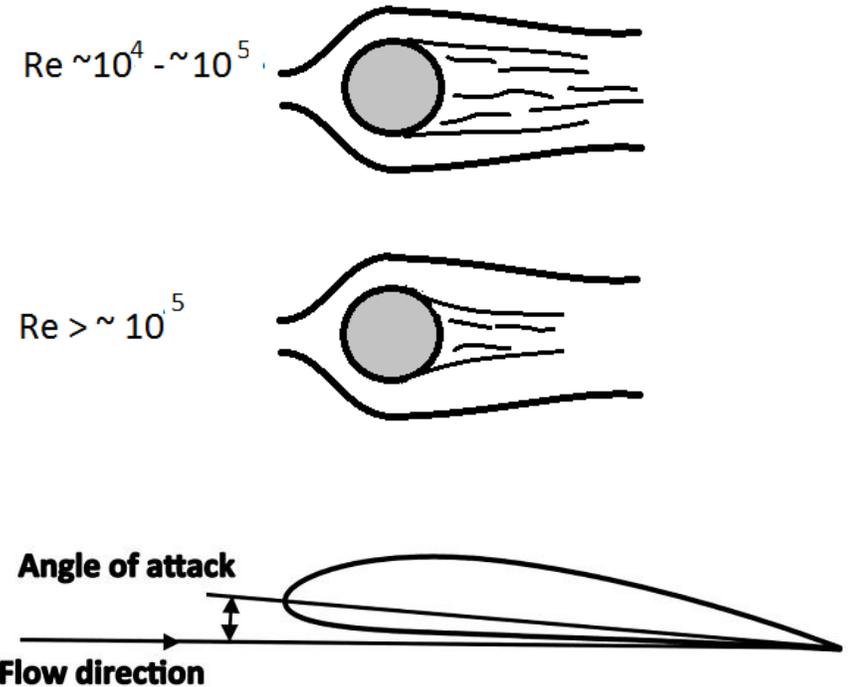


- 1. Resolution:** Consistent performance at varied scales
- 2. Data:** Benefiting from larger training data
- 3. Parameter:** Benefiting from more parameters



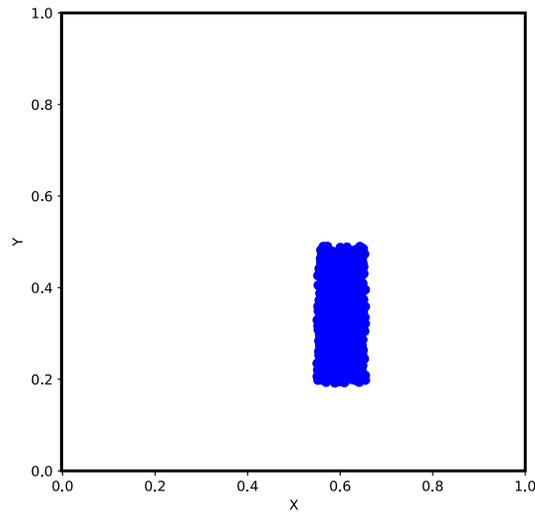
Pursuing PDE Foundation Models: Generalization

MODELS	OOD REYNOLDS		OOD ANGLES	
	$C_L \downarrow$	$\rho_L \uparrow$	$C_L \downarrow$	$\rho_L \uparrow$
SIMPLE MLP	0.6205	0.9578	0.4128	0.9572
GRAPHSAGE (2017)	0.4333	0.9707	<u>0.2538</u>	0.9894
POINTNET (2017)	0.3836	0.9806	0.4425	0.9784
GRAPH U-NET (2019)	0.4664	0.9645	0.3756	0.9816
MESHGRAPHNET (2021)	1.7718	0.7631	0.6525	0.8927
GNO (2020A)	0.4408	<u>0.9878</u>	0.3038	0.9884
GALERKIN (2021)	0.4615	0.9826	0.3814	0.9821
GNOT (2023)	<u>0.3268</u>	0.9865	0.3497	0.9868
GINO (2023A)	0.4180	0.9645	0.2583	<u>0.9923</u>
TRANSOLVER (OURS)	0.2996	0.9896	0.1500	0.9950

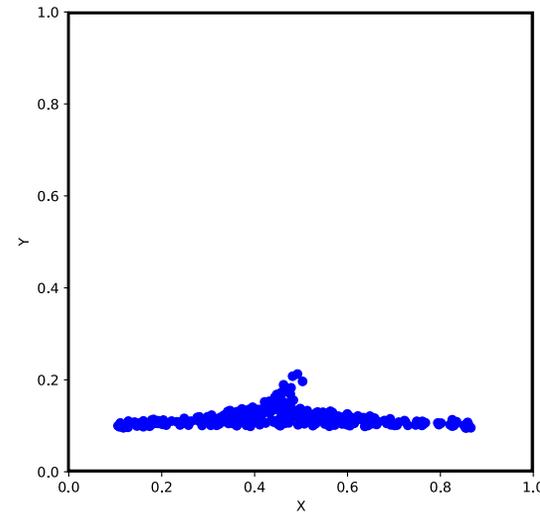


Transolver still performs best (**Spearman's correlation $\sim 99\%$**) in OOD settings

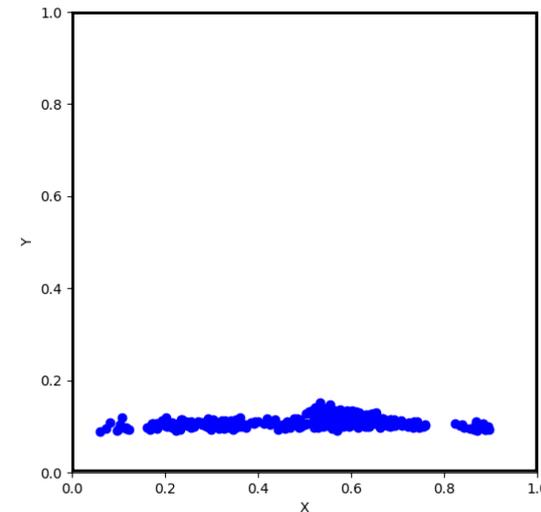
Pursuing PDE Foundation Models: Versatile



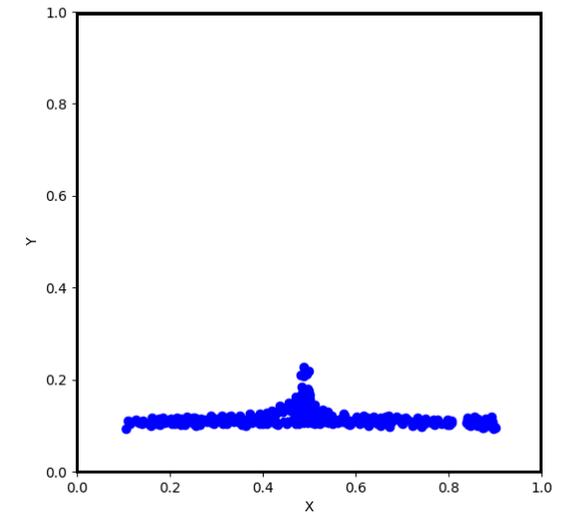
Initial State



Ground Truth (400th step)



GNN



GNN + Transolver

MODEL	MSE ↓
GNN (SANCHEZ-GONZALEZ ET AL., 2020)	0.0182
GNN + TRANSOLVER (OURS)	0.0069
RELATIVE PROMOTION	62.1%

Transolver can also be extended to

Lagrangian Settings
(Ever-changing geometrics)

Open Source

thuml / Transolver

Code Issues Pull requests Actions Projects Security Insights Settings

Transolver Public

main 1 Branch 0 Tags

Go to file Add file Code

wuhaixu2016 Update exp_elas.py 9e0addd · 2 days ago 7 Commits

Airfoil-Design-AirFRANS	Update requirements.txt	3 days ago
Car-Design-ShapeNetCar	update vis	3 days ago
PDE-Solving-StandardBenchmark	Update exp_elas.py	2 days ago
pic	init code	3 days ago
.gitignore	Initial commit	last week
LICENSE	Initial commit	last week
Physics_Attention.py	init code	3 days ago
README.md	init code	3 days ago

README MIT license

Transolver (ICML 2024)

Transolver: A Fast Transformer Solver for PDEs on General Geometries [\[paper\]](#)

In real-world applications, PDEs are typically discretized into large-scale meshes with complex geometries. To capture intricate physical correlations hidden under multifarious meshes, we propose the Transolver with the following features:

- Going beyond previous work, Transolver **calculates attention among learned physical states** instead of mesh points, which empowers the model with **endogenetic geometry-general capability**.
- Transolver achieves **22% error reduction over previous SOTA in six standard benchmarks** and excels in **large-scale industrial simulations**, including car and airfoil designs.

About code release of "Transolver: A Fast Transformer Solver for PDEs on General Geometries", ICML 2024. <https://arxiv.org/abs/2402.02366>

Readme MIT license Activity Custom properties 1 star 3 watching 0 forks Report repository

Releases No releases published [Create a new release](#)

Packages No packages published [Publish your first package](#)

Languages

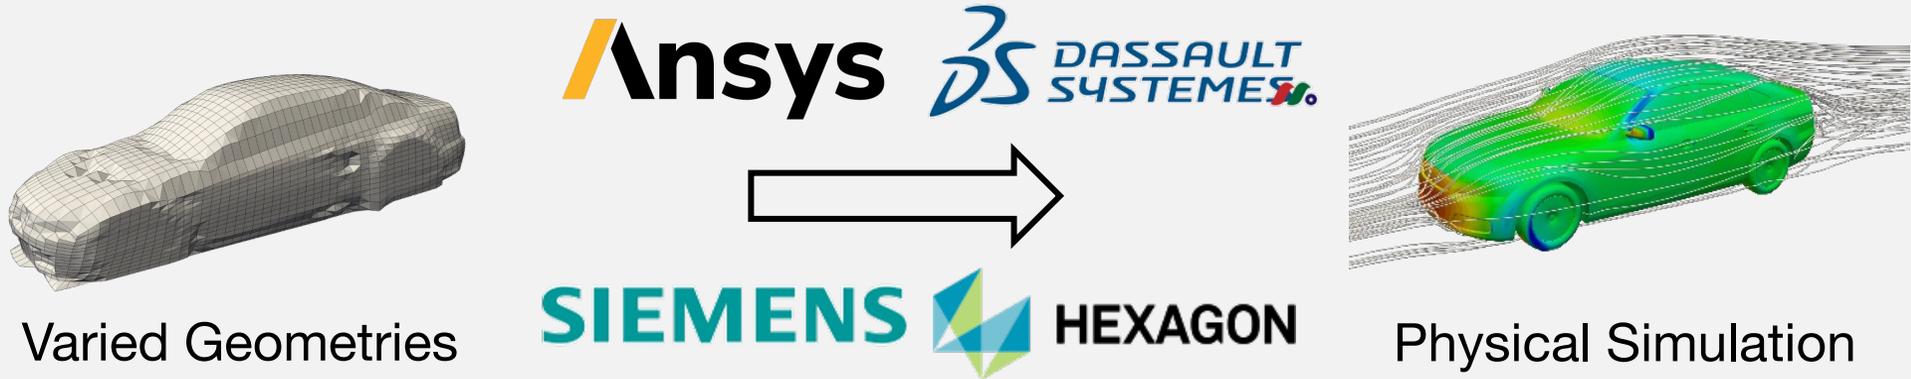
- Python 97.9%
- Jupyter Notebook 1.3%
- Shell 0.8%

Suggested workflows Based on your tech stack

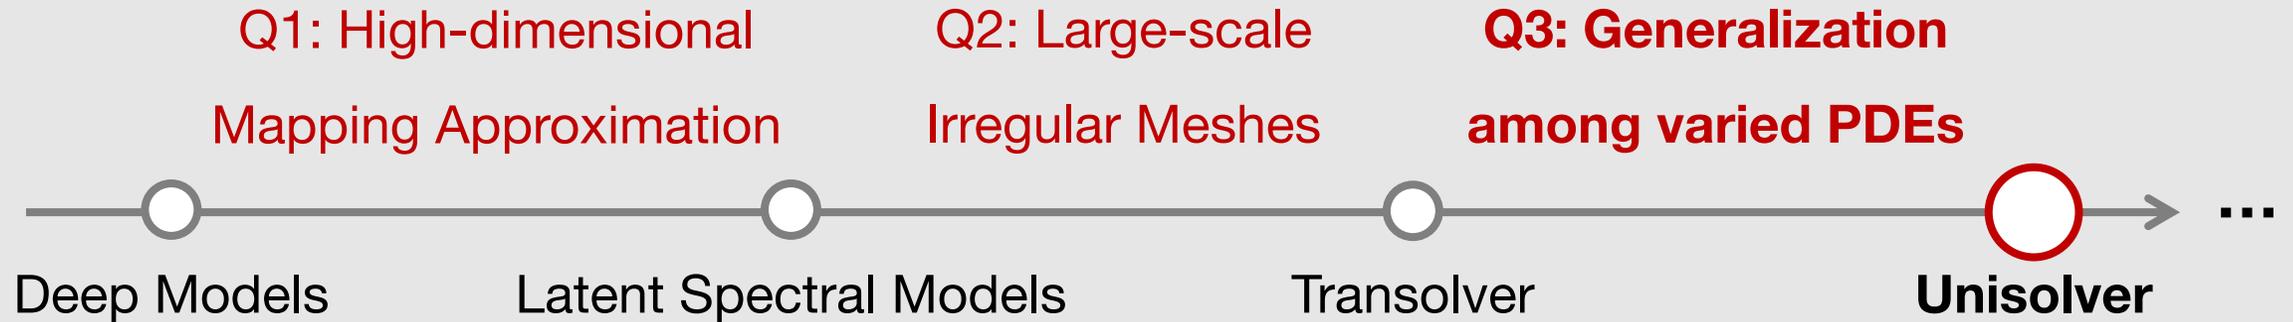
Code is available at <https://github.com/thuml/Transolver>

A Roadmap to Practical Neural PDE Solvers

Industrial simulation with CAE



Neural PDE Solver (Our work)





Unisolver: PDE-Conditional Transformers Are Universal PDE Solvers

Hang Zhou*, Yuezhou Ma*, Haixu Wu[✉], Haowen Wang, Mingsheng Long[✉]
School of Software, BNRist, Tsinghua University, China



Hang Zhou



Yuezhou Ma



Haixu Wu

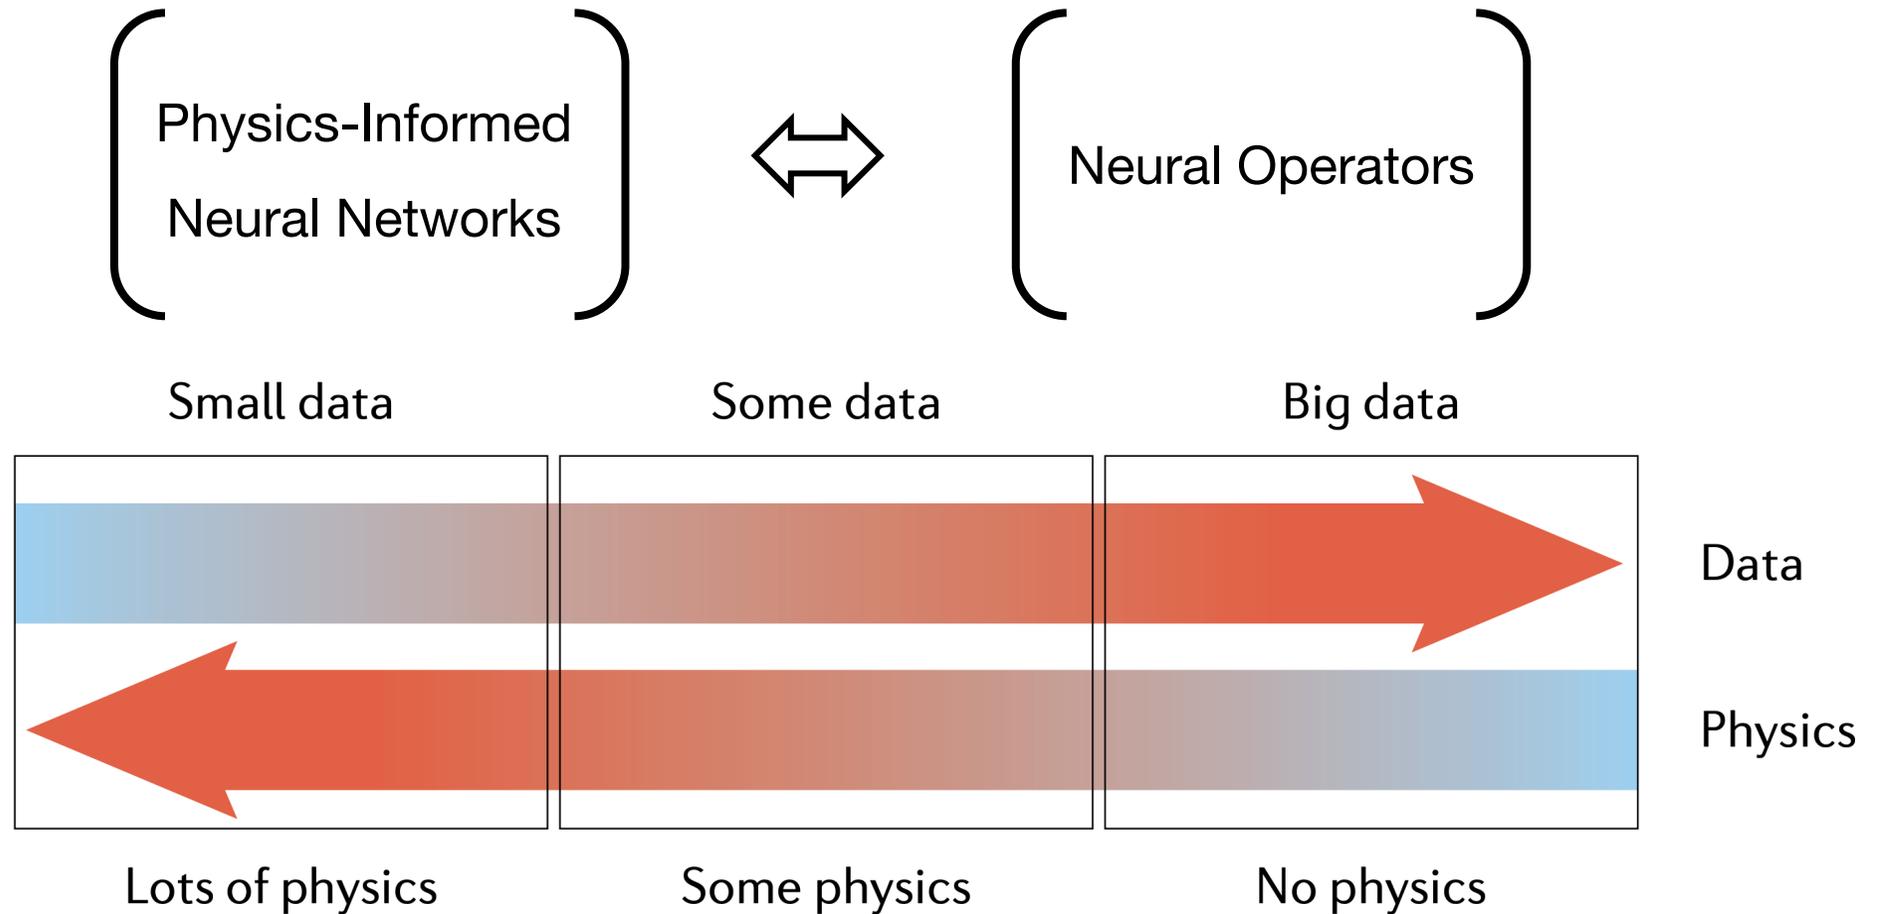


Haowen Wang



Mingsheng Long

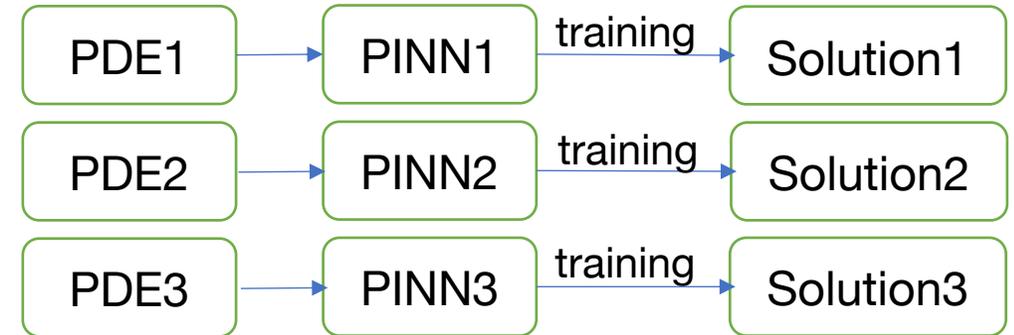
Machine Learning for PDEs



Generalizability of Neural PDE Solvers

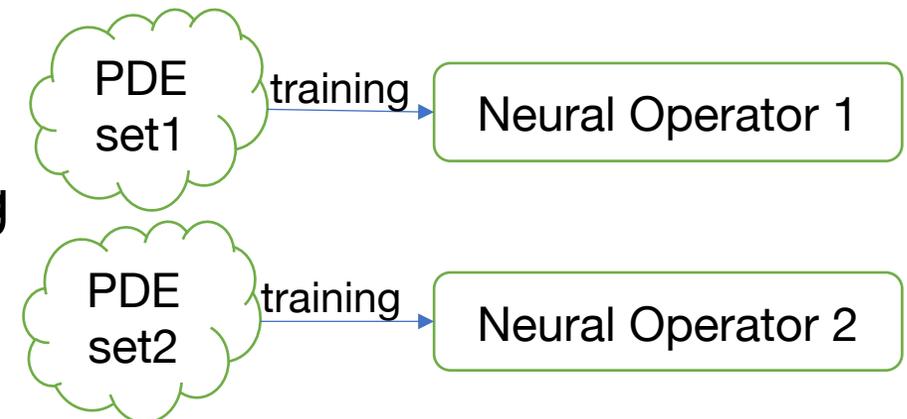
➤ PINNs

1. Formalize the specific PDE equations as objective functions during training
2. Struggle to generalize to unseen PDEs
3. Re-training required to solve new PDEs

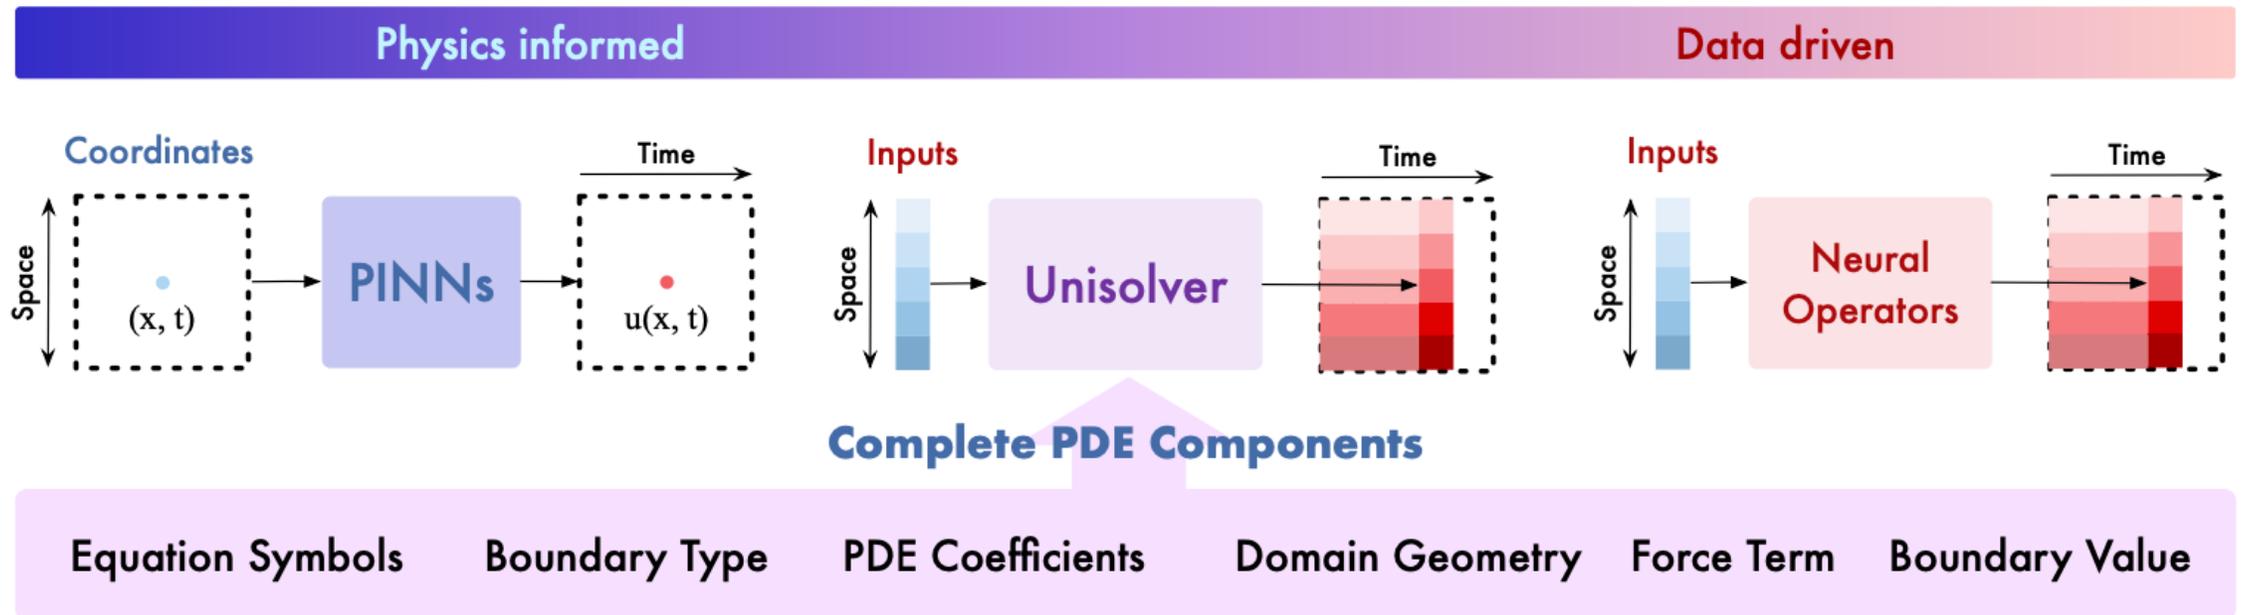


➤ Neural operators

1. learning from pre-computed simulation data
2. Cannot efficiently adapt to PDEs with varying components
3. Computational costly and data-demanding



Unisolver: A Unification of Two Paradigms



In addition to simulated data, Unisolver also defines and utilizes **a complete set of PDE components.**

Complete PDE Components

Motivating example: vibrating string equation

$$\partial_{tt}u - a^2\partial_{xx}u = f(x, t), \quad (x, t) \in (0, L) \times (0, T), \quad (1a)$$

$$u(0, t) = 0, \quad u(L, t) = 0, \quad t \in (0, T], \quad (1b)$$

$$u(x, 0) = \phi(x), \quad \partial_t u(x, 0) = \psi(x), \quad x \in [0, L]. \quad (1c)$$

- The **coefficient** a represents physical quantity such as tension, linear density
- f represents the **external force** driving the vibrations of the string
- Equation (1b) sets **boundary conditions** at endpoints
- Equation (1c) specifies **initial conditions**
- The **domain geometry** spans the range $[0, L] \times [0, T]$

Complete PDE Components

Motivating example: vibrating string equation

$$\partial_{tt}u - a^2\partial_{xx}u = f(x, t), \quad (x, t) \in (0, L) \times (0, T), \quad (1a)$$

$$u(0, t) = 0, \quad u(L, t) = 0, \quad t \in (0, T], \quad (1b)$$

$$u(x, 0) = \phi(x), \quad \partial_t u(x, 0) = \psi(x), \quad x \in [0, L]. \quad (1c)$$

- The analytical solution of the above equations is:

$$u(x, t) = \frac{1}{2} \underbrace{(\Phi(x + at) + \Phi(x - at))}_{\text{Initial position}} + \frac{1}{2a} \underbrace{\int_{x-at}^{x+at} \Psi(\xi) d\xi}_{\text{Initial velocity}} + \frac{1}{2a} \underbrace{\int_0^t d\tau \int_{x-a(t-\tau)}^{x+a(t-\tau)} f(\xi, \tau) d\xi}_{\text{Geometry Force}}$$

- The PDE is solved under complex interactions between equation components
- The impact of the **external force** is imposed **point-wisely**
- The **coefficient** exerts a **consistent** influence over the domain

Complete PDE Components

Categorization of PDE components

Category	Component	Description
Domain-wise components	Equation formulation	The symbolic expression of the PDE
	Equation coefficient	Coefficients in the PDE equation
	Boundary condition type	Type of boundary condition (e.g., Dirichlet, Neumann)
Point-wise components	External force	Forces acting at specific points
	Domain geometry	The shape and size of the domain
	Boundary value function	Value functions at the domain boundaries

- Category PDE components into **domain-** and **point-wise** components:
- Here the equation formulation refers to the **symbolic expression of PDEs**, which can be encoded by **Large Language Models**

Universal Components Embedding

□ Embedding of equation formulation

- ✓ Utilizing a LLM to embed the symbolic expression of the PDE
- ✓ The symbolic expression is represented by the LaTeX code

Prompt: "`\partial_{tt} u - a^2\partial_{xx} u = f(x,t)`"

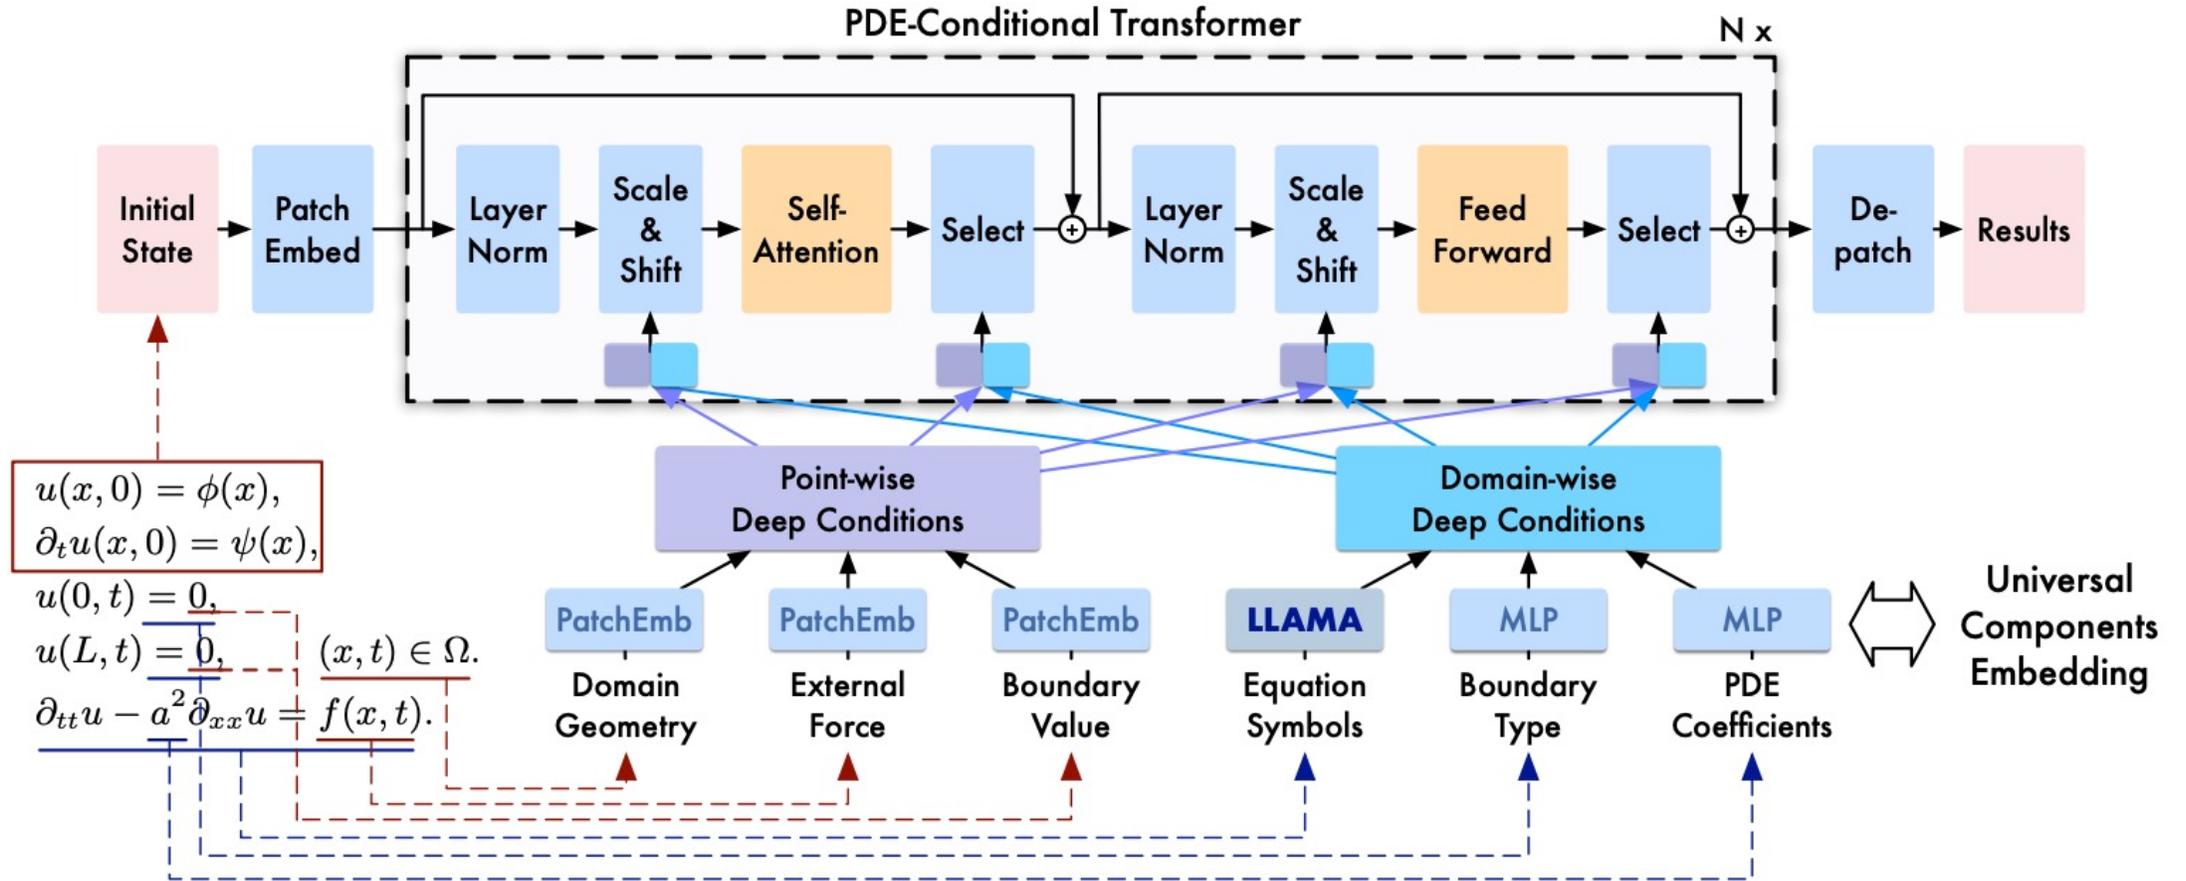
□ Embedding of other components

- ✓ **Domain-wise** components are embedded by a 2-layer MLP
- ✓ **Point-wise** components are patchified and embedded into tokens

□ Deep condition consolidation

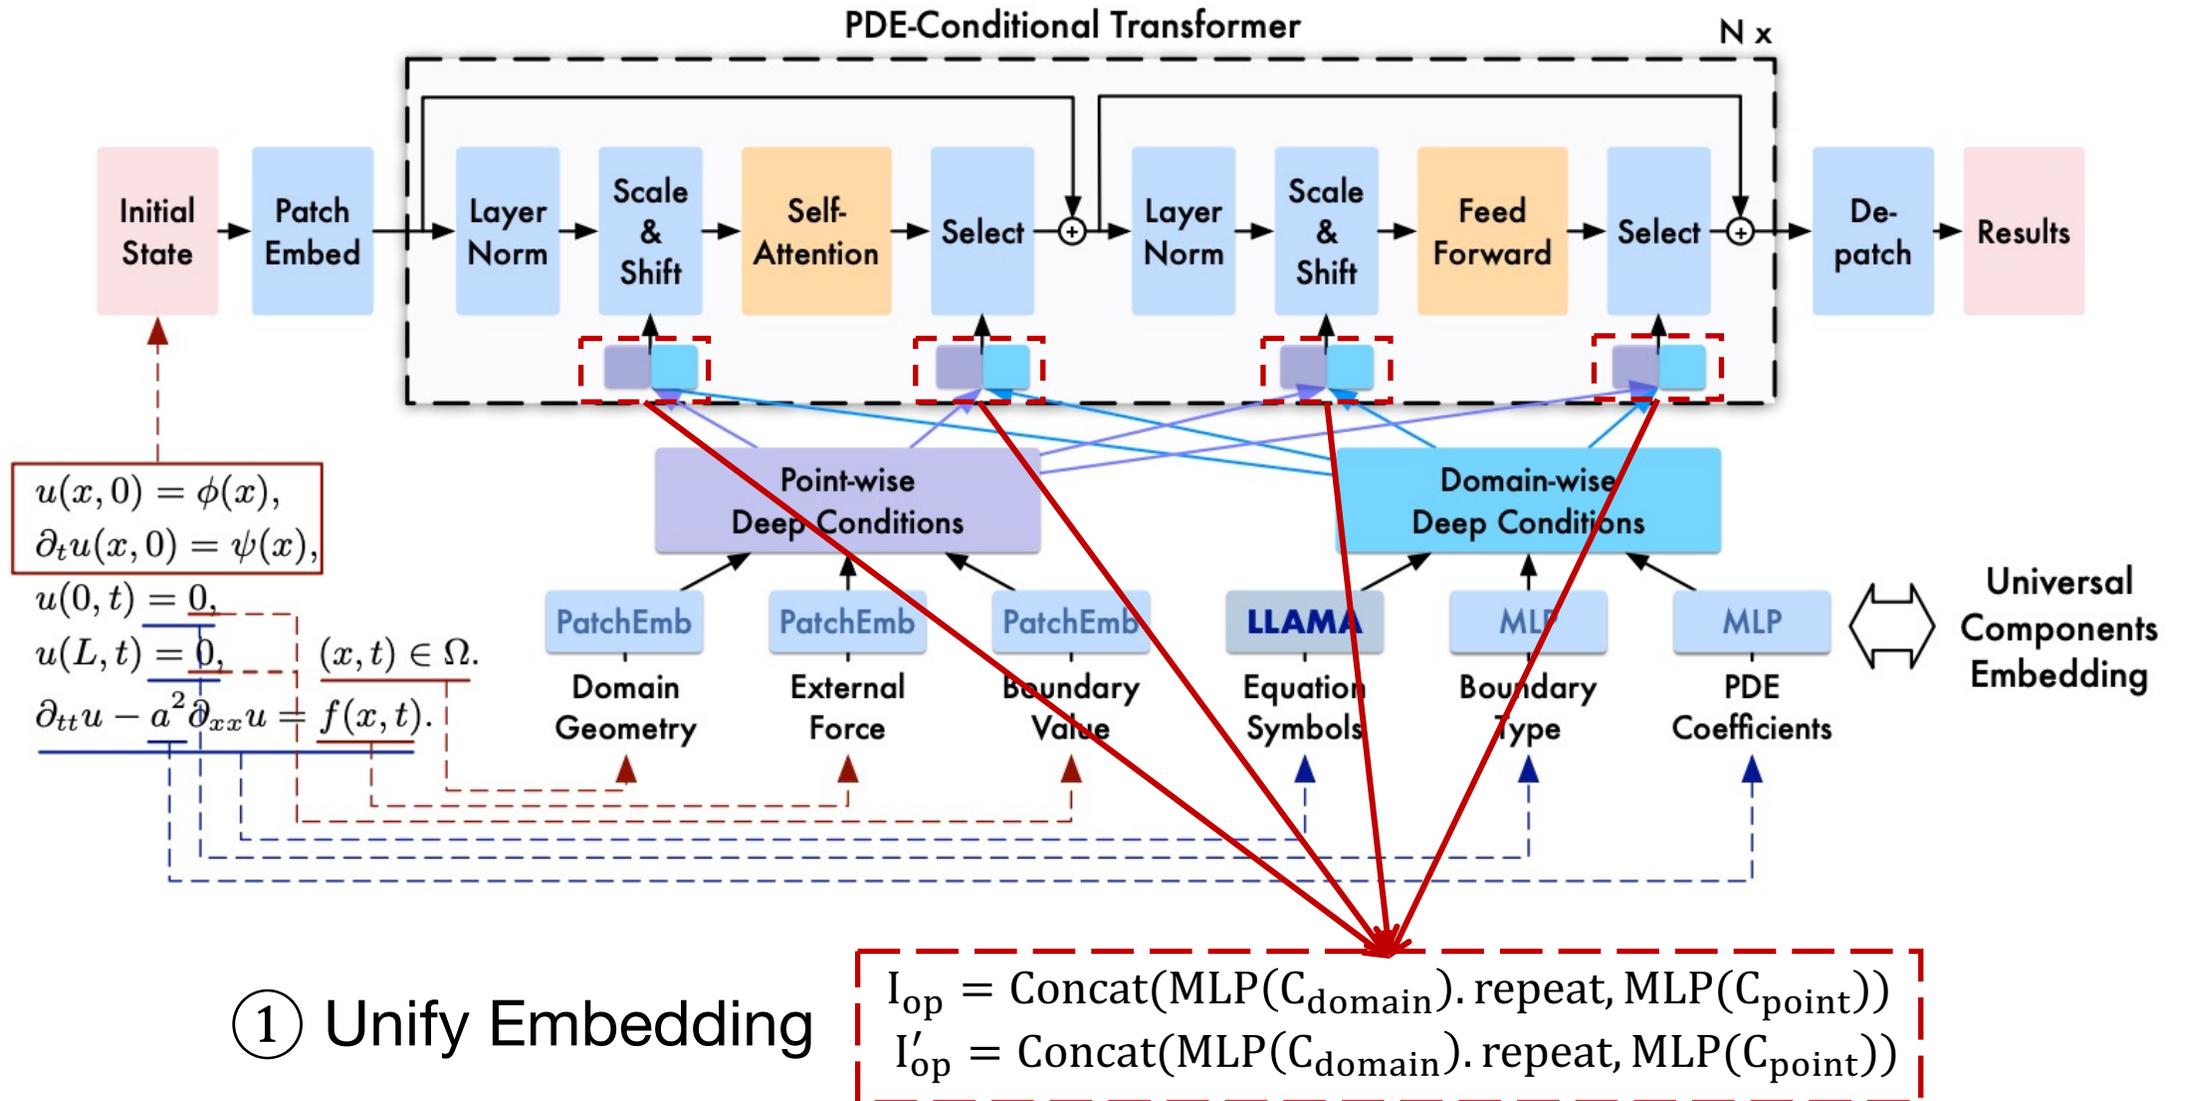
- ✓ Embedded conditions of each categories are aggregated together
- ✓ Make for easy adaptation to novel PDE components in downstream tasks

PDE-Conditional Transformer

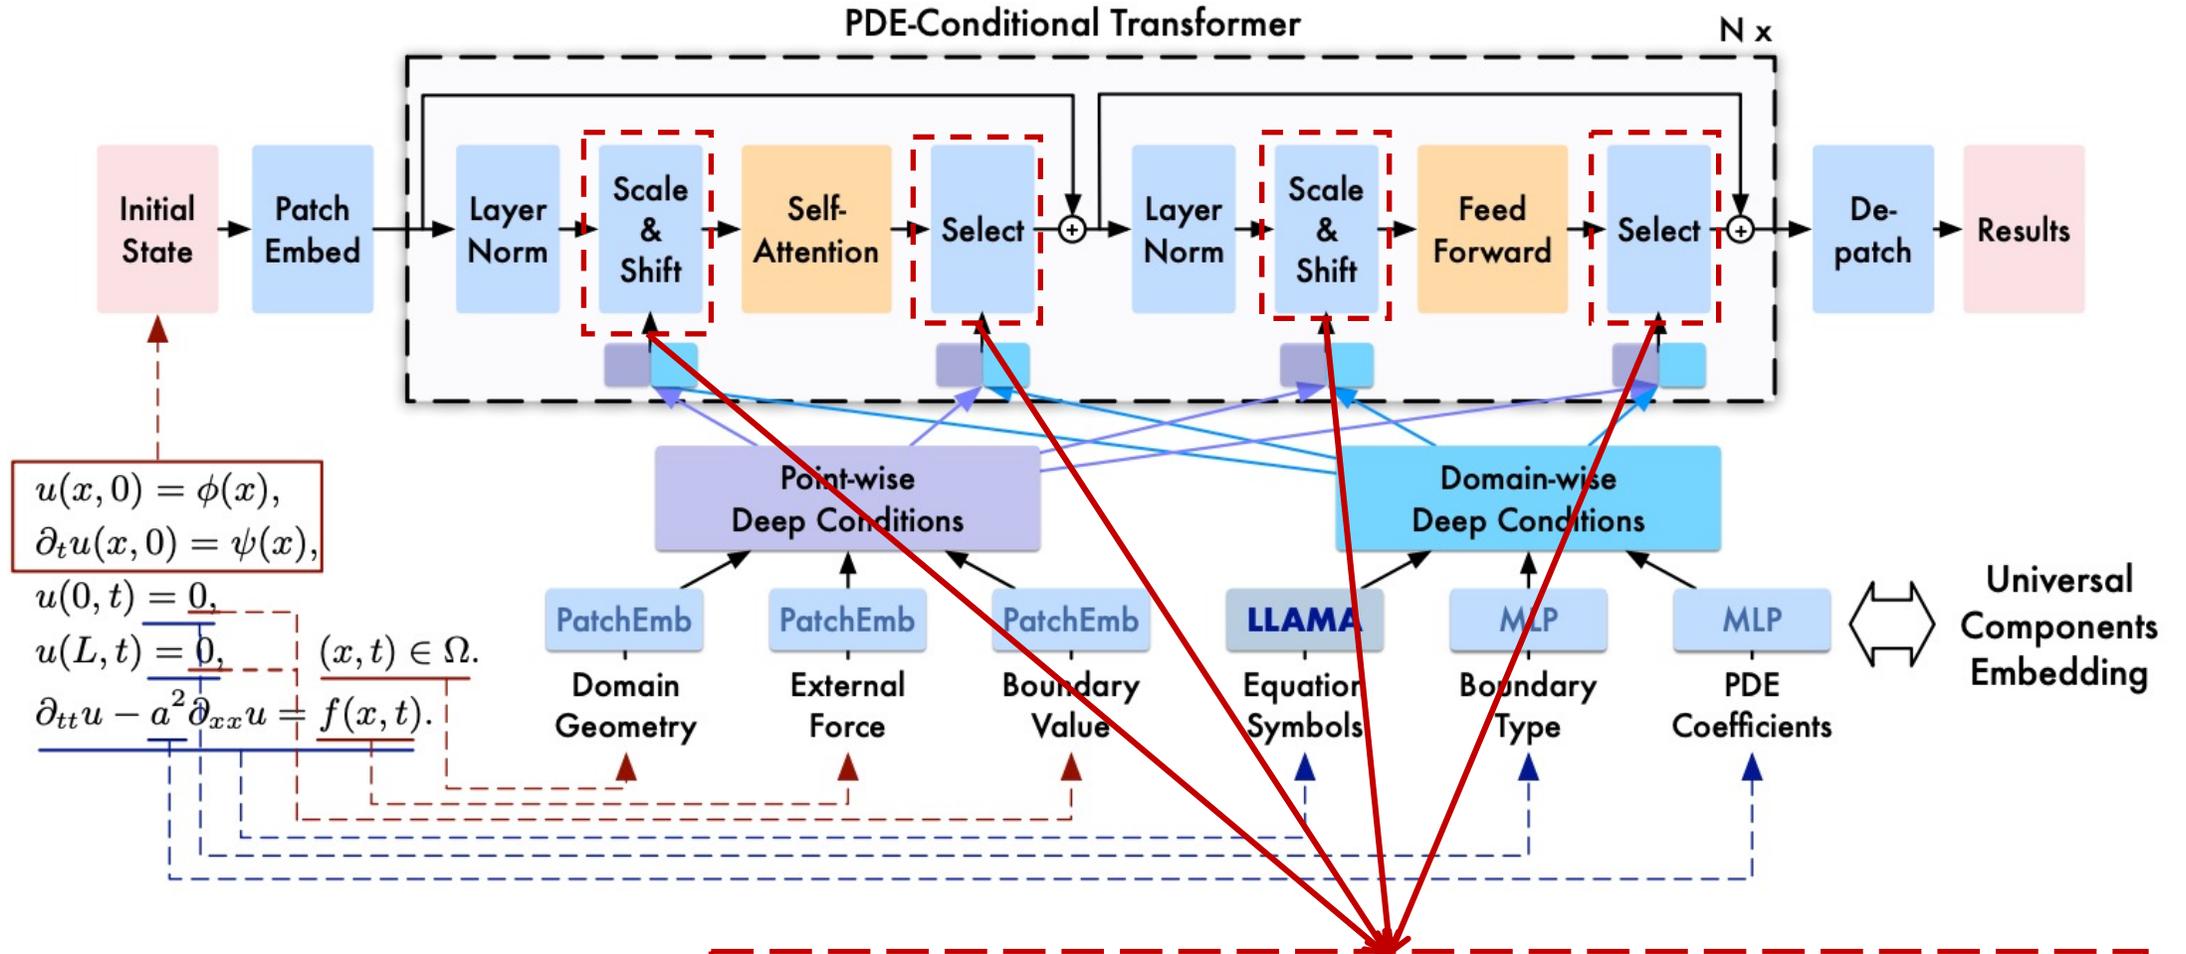


- ① Unify Embedding ② Condition Aggregation

PDE-Conditional Transformer



PDE-Conditional Transformer



② Condition Aggregation

$$X^{n-1/2} = I_{\text{select}} \odot \text{Attn}(I_{\text{scale}} \odot \text{LayerNorm}(X^{n-1}) + I_{\text{shift}}) + X^{n-1}$$

$$X^n = I'_{\text{select}} \odot \text{FFN}(I'_{\text{scale}} \odot \text{LayerNorm}(X^{n-1/2}) + I'_{\text{shift}}) + X^{n-1/2}$$

Experiments

Benchmarks	#Dim	#Resolution	# Samples	#Size	Symbols	Coefficient	Force	Geometry	Boundary
HeterNS	2D+Time	(64,64,10)	15k	4.6 GB	×	✓	✓	×	×
PDEformer [45]	1D+Time	(256,100)	3M	300 GB	✓	✓	✓	×	✓
DPOT [10]	2D+Time	(128,128,10)	74.1k	384 GB	✓	✓	✓	✓	✓

- HeterNS contains multiple **viscosity coefficients** and **external force**
- PDEformer proposes a large-scale dataset with 3M samples of 1D PDEs, including multiple equation **coefficients**, **external force** and **boundary conditions**
- DPOT collects 12 datasets from FNO, PDEBench, PDEArena and CFDBench, with PDEs varying in **coefficients**, **external force**, **geometries** and **boundary conditions**

Heterogeneous 2D Navier-Stokes Equation

- Generalize to unseen coefficients

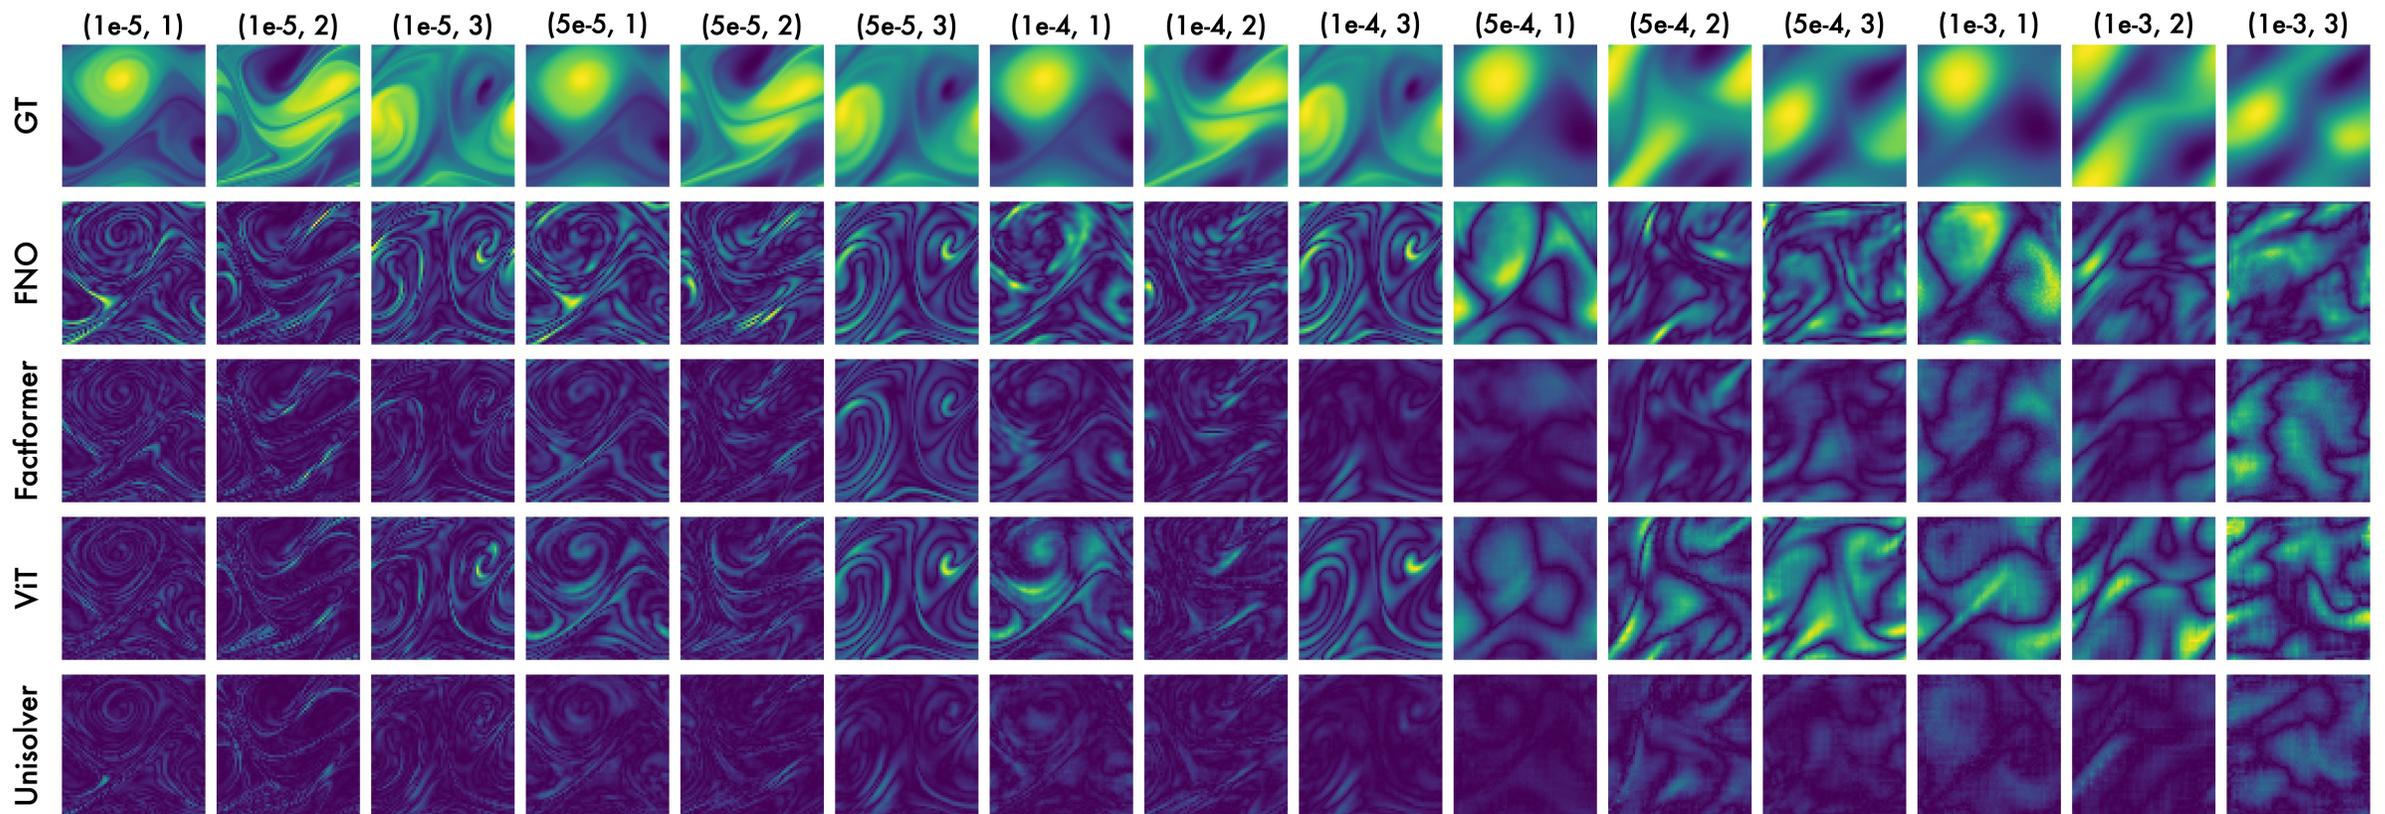
HeterNS	Viscosity ν	8e-6	1e-5	3e-5	5e-5	8e-5	1e-4	3e-4	5e-4	8e-4	1e-3	2e-3
	Params	OOD	ID	OOD								
FNO	4.7M	0.0702	0.0669	0.0373	0.0225	0.0141	0.0114	0.0088	0.0031	0.0084	0.0011	0.2057
Factformer	9.4M	0.0489	0.0438	0.0489	0.0128	0.0297	0.0064	0.1386	0.0018	0.0631	0.0010	0.3207
ViT	4.8M	0.0458	0.0432	0.0353	0.0206	0.0119	0.0098	0.0100	0.0031	0.0174	0.0015	0.1878
Unisolver	4.1M	0.0336	0.0321	0.0178	0.0094	0.0064	0.0051	0.0066	0.0015	0.0096	0.0008	0.1504
Promotion	-	26.6%	25.7%	49.6%	26.6%	46.2%	20.3%	34.0%	16.7%	-	20.0%	19.9%

- Generalize to unseen external force

HeterNS	Force ω	0.5	1	1.5	2	2.5	3	3.5
	Params	OOD	ID	OOD	ID	OOD	ID	OOD
FNO	4.7M	1.110	0.0640	0.1742	0.0661	0.1449	0.1623	0.2974
Factformer	9.4M	0.9998	0.0326	0.1110	0.0438	0.1243	0.0803	0.2257
ViT	4.8M	0.7900	0.0348	0.1412	0.0432	0.1240	0.1000	0.2080
Unisolver	4.1M	0.0980	0.0244	0.0770	0.0321	0.0720	0.0720	0.1740
Promotion	-	87.6%	25.2%	30.6%	25.7%	41.9%	10.3%	16.4%

Heterogeneous 2D Navier-Stokes Equation

- All showcases generated with the same initial condition but with varied coefficients. **Different viscosities presents quite different dynamics.**

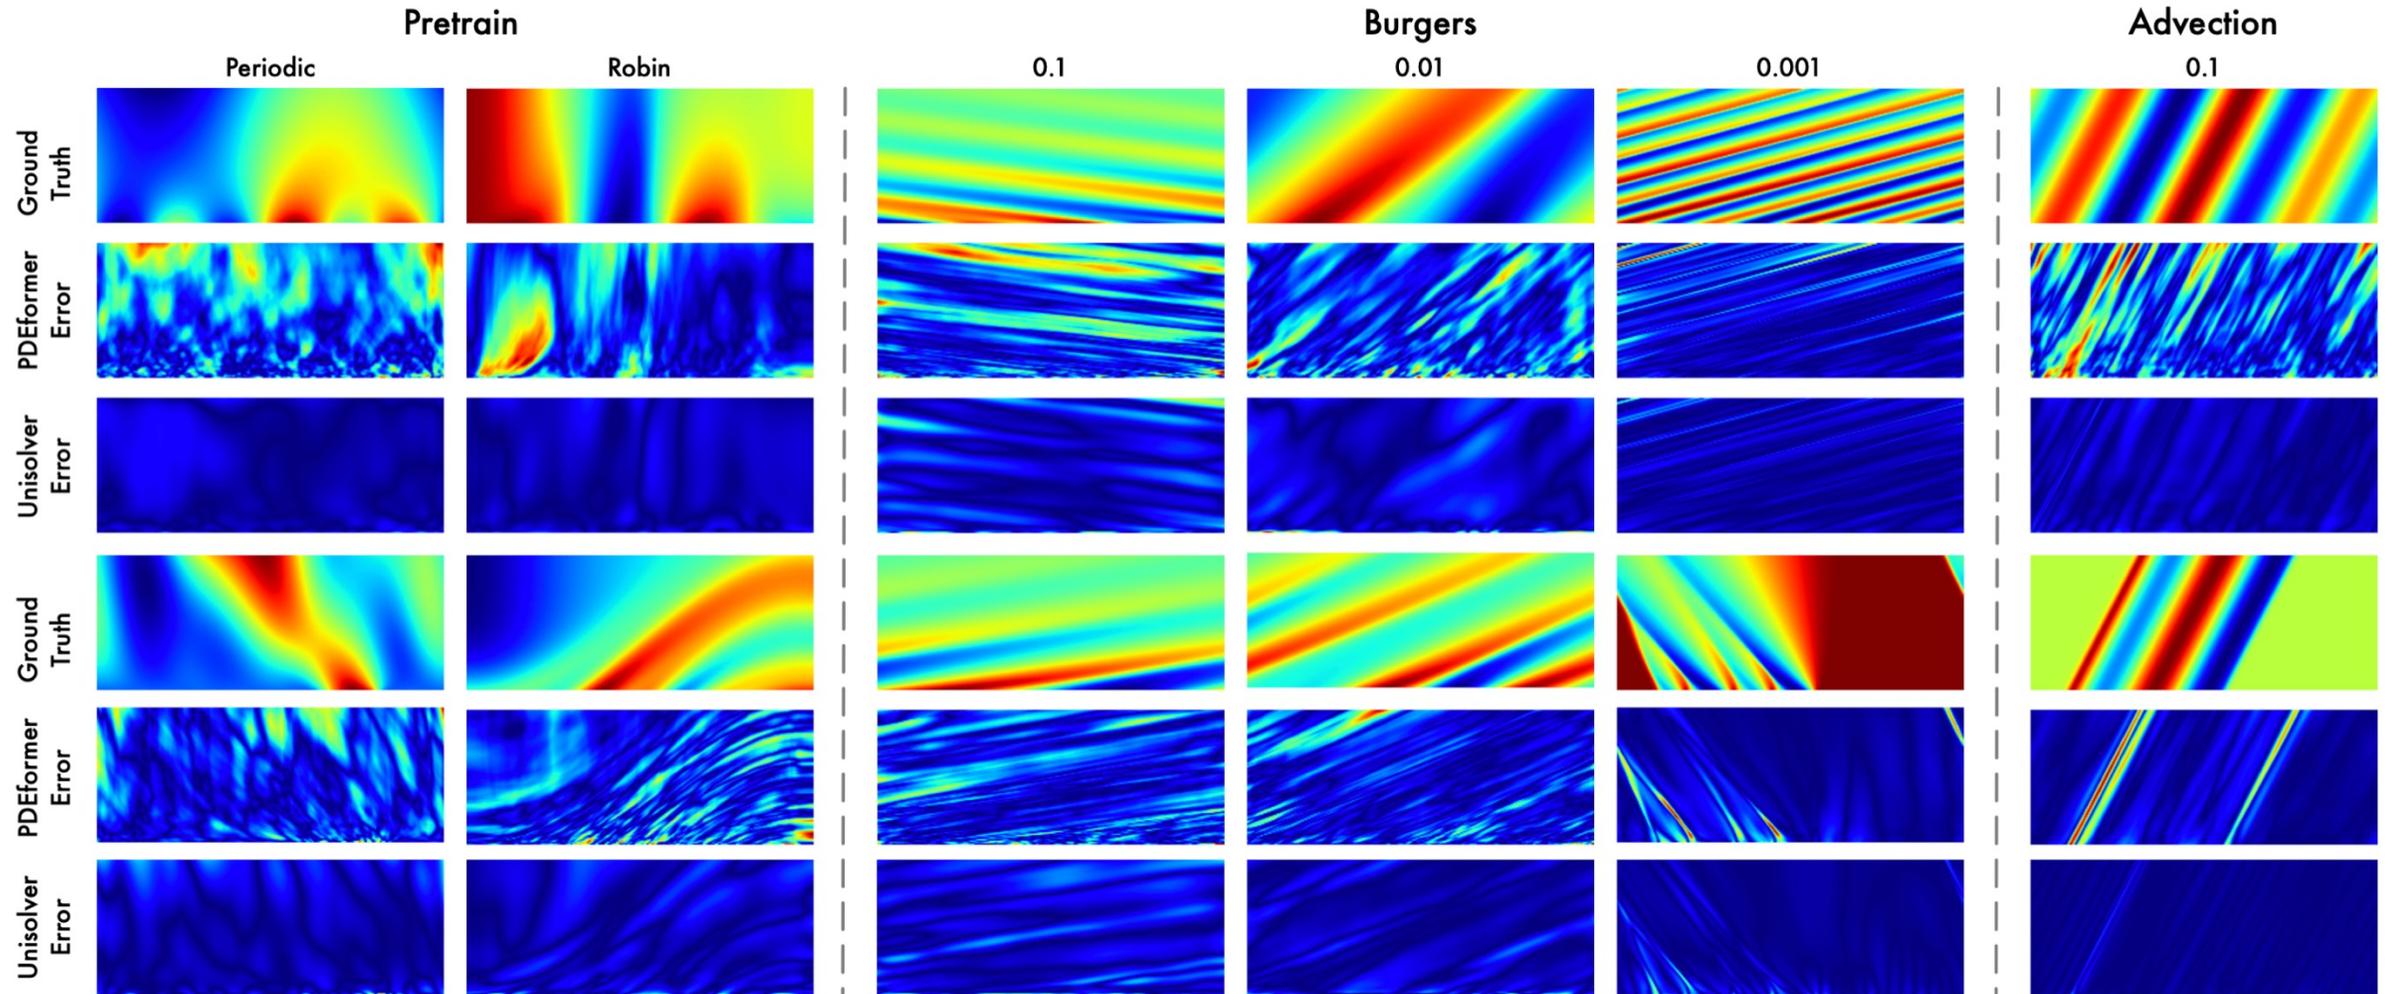


1D Time-dependent PDEs proposed by PDEformer

- Models are pretrained on a dataset with **3 million 1D PDEs** with varied **coefficients, external force, boundary conditions and equation symbols**
- Then tested on **OOD downstream PDE datasets** generated by PDEBench, including the Burgers equation and the advection equation

1D Time-dependent PDEs	Tasks	Pretrain		Burgers			Advection
	Params	Periodic	Robin	$\nu = 0.1$	$\nu = 0.01$	$\nu = 0.001$	$\beta = 0.1$
PDEformer-L	22M	0.0211	0.0238	0.00744	0.0144	0.0393	0.0178
Unisolver	19M	0.0107	0.0108	0.00513	0.00995	0.0299	0.0138
Promotion	-	49.3%	54.6%	31.0%	30.9%	23.9%	22.5%
PDEformer-L (FT-100)	22M	-	-	0.00364	0.0112	0.0331	0.00975
Unisolver (FT-100)	19M	-	-	0.00105	0.00474	0.0170	0.00420
Promotion	-	-	-	71.2%	57.7%	48.6%	59.6%

Showcases

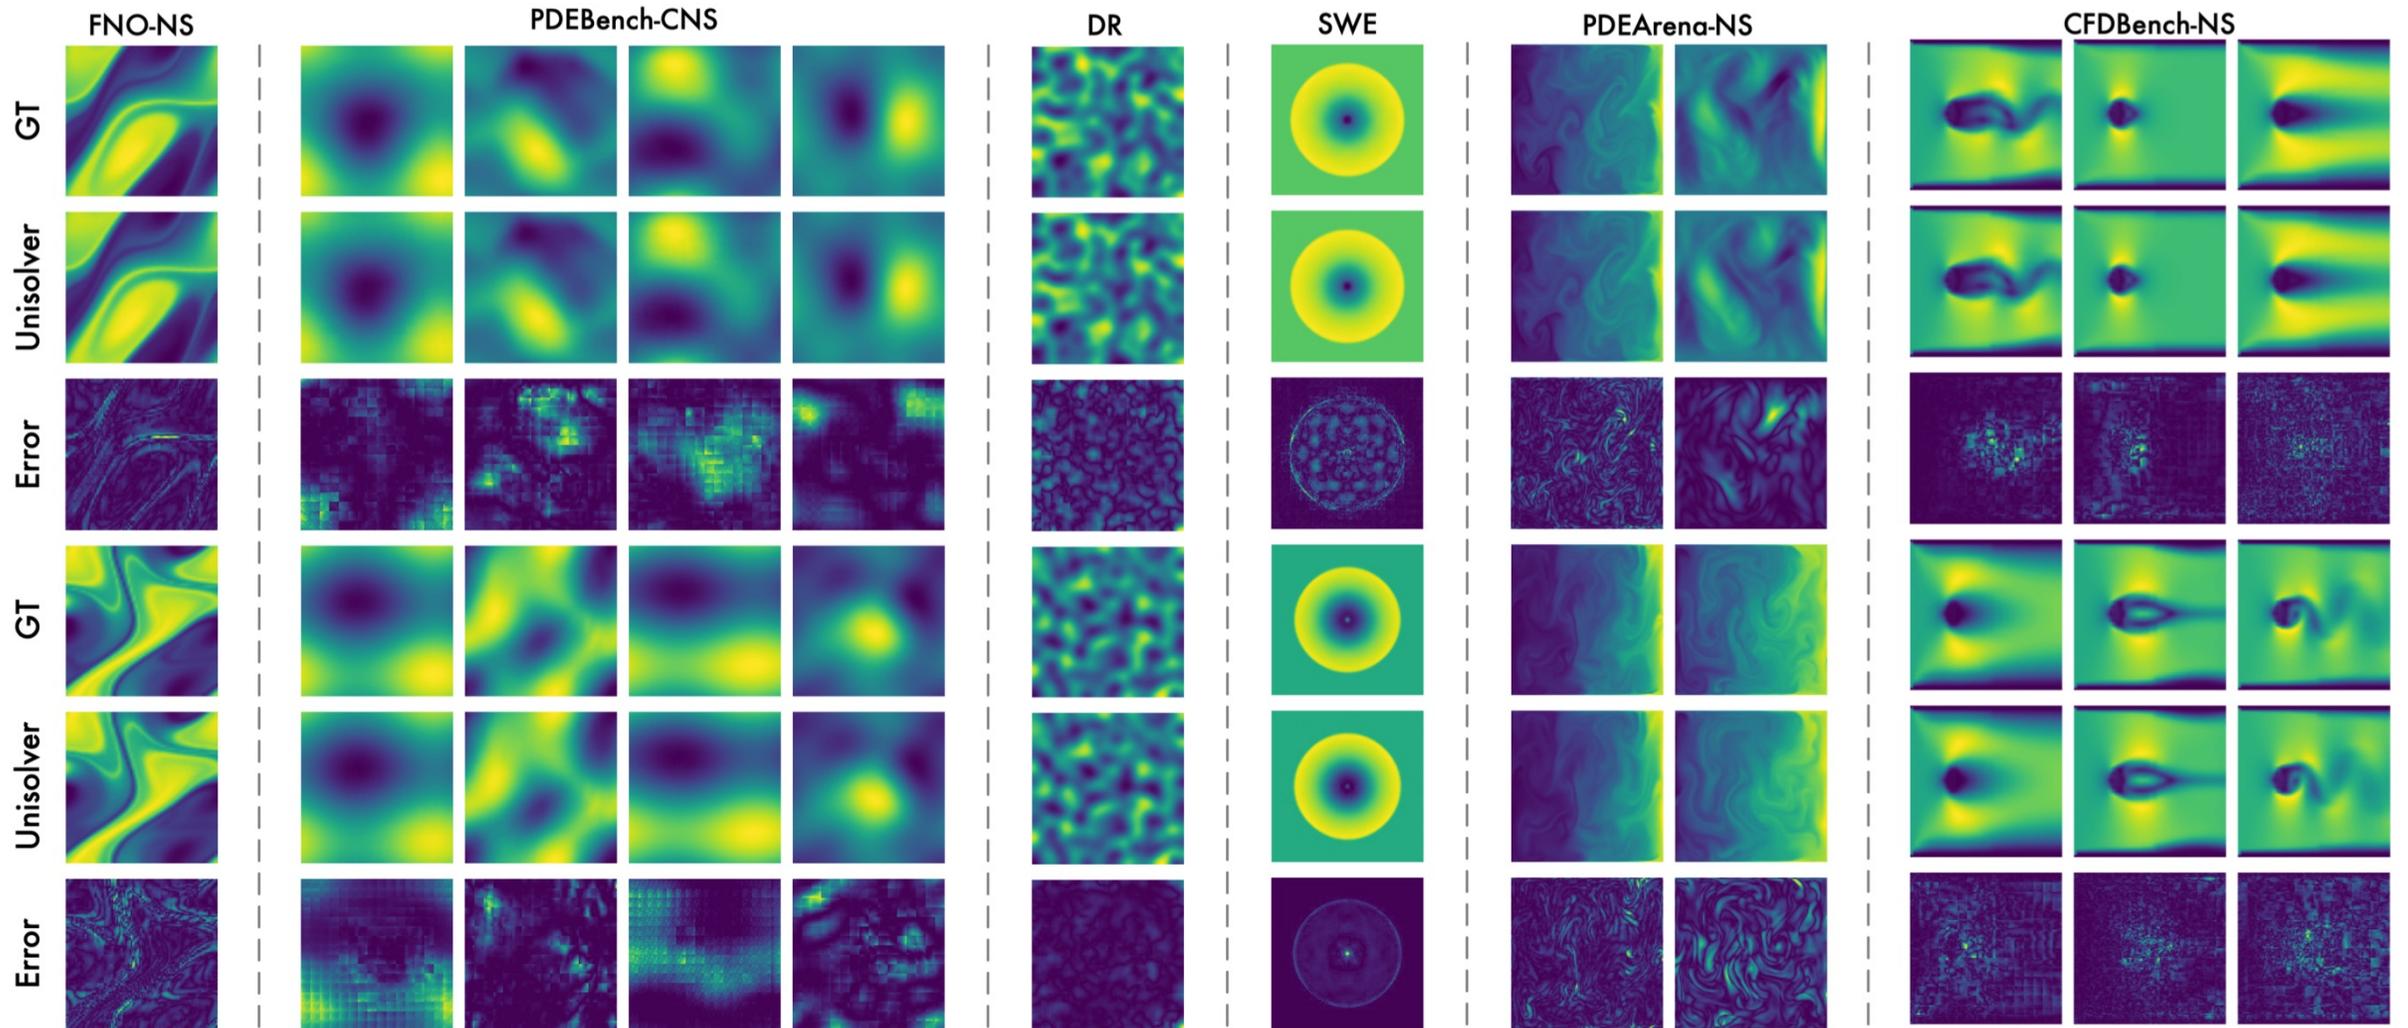


2D Mixed PDEs proposed by DPOT

The models are pretrained on 12 datasets collected by DPOT, with varied coefficients, external force, boundary conditions and geometries

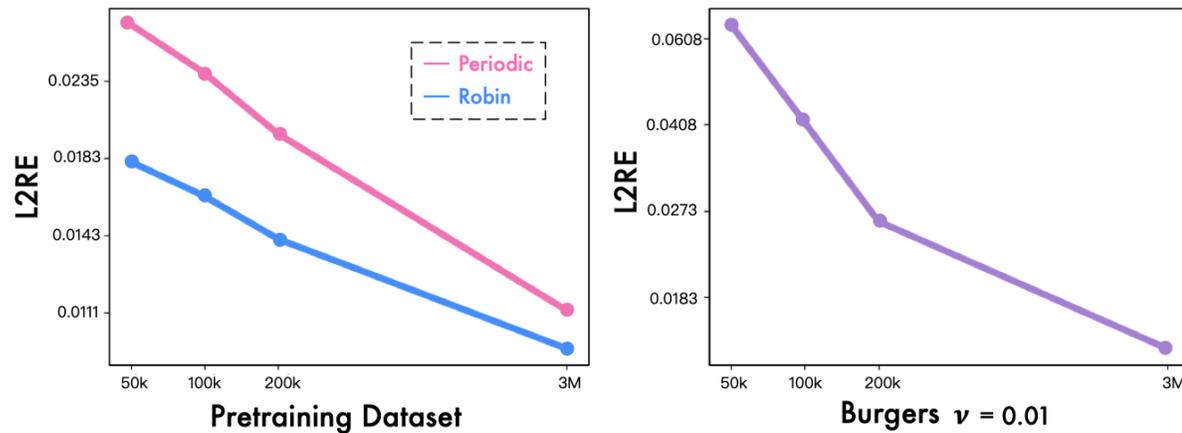
2D Mixed PDEs	Equations Params	FNO-NS- ν			PDEBench-CNS-(M, ζ)				DR	SWE	PDEArena-NS		CFDBench-NS Geometry	Mean
		1e-5	1e-4	1e-3	(1, 0.1)	(1, 0.01)	(0.1, 0.1)	(0.1, 0.01)			-	-		
DPOT-S	30M	5.53	4.42	1.31	1.53	3.37	1.19	1.87	3.79	0.66	9.91	31.6	0.70	5.50
Unisolver	33M	4.17	3.36	0.61	1.23	2.89	1.01	1.59	4.39	0.45	6.87	27.4	0.54	4.54
Promotion	-	24.6%	24.0%	53.4%	19.6%	14.2%	15.1%	15.0%	-	31.8%	30.7%	13.3%	22.9%	17.5%
DPOT-S-FT	30M	4.49	3.42	0.68	1.52	2.11	1.50	1.51	1.71	0.22	8.92	29.0	0.44	4.63
Unisolver-FT	33M	3.82	2.79	0.31	0.95	1.99	1.01	1.34	1.37	0.20	6.67	26.8	0.47	3.98
Promotion	-	14.9%	18.4%	54.4%	37.5%	5.69%	32.6%	11.3%	19.8%	9.1%	25.2%	7.6%	-	14.0%

Showcases

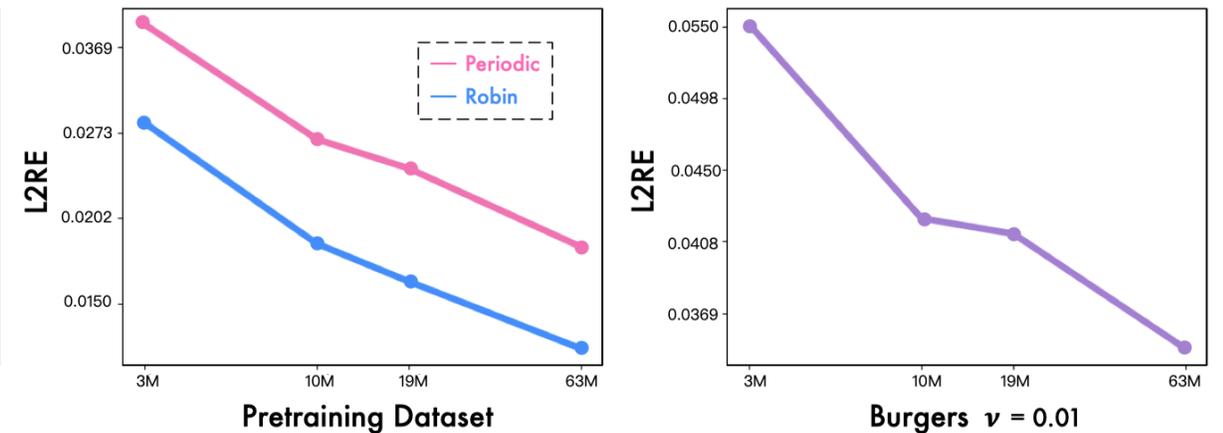


Scalability

(a) Data Scalability (Samples)



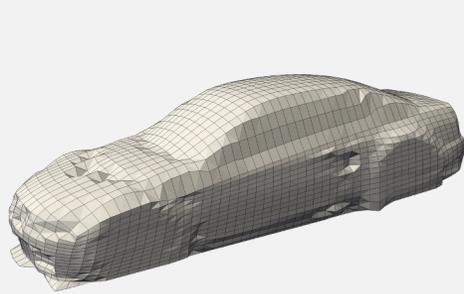
(b) Model Scalability (Parameters)



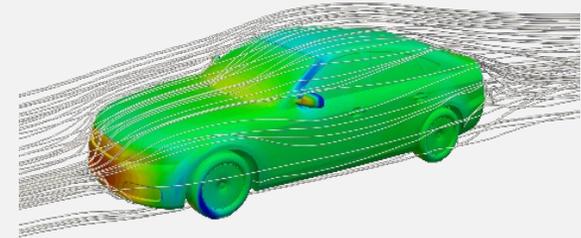
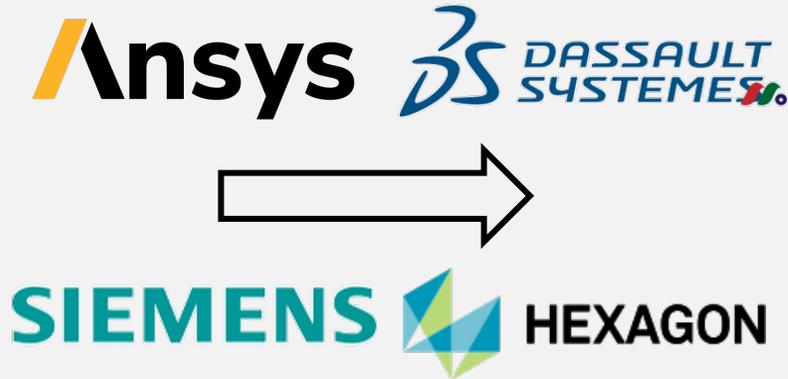
We progressively increase the **training data by 60 times** and **the model parameters by 21 times**, plotting the Relative L2 error on a log-log scale

A Roadmap to Practical Neural PDE Solvers

Industrial simulation with CAE



Varied Geometries



Physical Simulation

Neural PDE Solver (Our work)

Q1: High-dimensional
Mapping Approximation

Q2: Large-scale
Irregular Meshes

Q3: Generalization
among varied PDEs



A Roadmap to Practical Neural PDE Solvers

Mingsheng Long

School of Software, Tsinghua University

July 2024



清華大學
Tsinghua University

