



ICML | 2024

The Forty-first International Conference on Machine Learning



---

# Transolver: A Fast Transformer Solver for PDEs on General Geometries

---

Haixu Wu<sup>1</sup> Huakun Luo<sup>1</sup> Haowen Wang<sup>1</sup> Jianmin Wang<sup>1</sup> Mingsheng Long<sup>1</sup>



Haixu Wu



Huakun Luo



Haowen Wang



Jianmin Wang



Mingsheng Long


# 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

# Solving PDEs

## Classic Numerical Methods




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

**Huge computation costs**

# Solving PDEs


## Classic Numerical Methods



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

**Huge computation costs**

## Neural PDE Solver




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




**An efficient surrogate tool**





# Solving PDEs: Discretization




Car



Airplane




# Challenges in Practical Industrial Design



Task: Estimate the drag coefficient of a given shape:

**Surrounding Wind & Surface Pressure**

# Challenges in Practical Industrial Design




Task: 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*

# 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*

# Previous Work: Geometry-General Neural Operators



## (1) GNN as Operators

*GNO, GINO, etc*

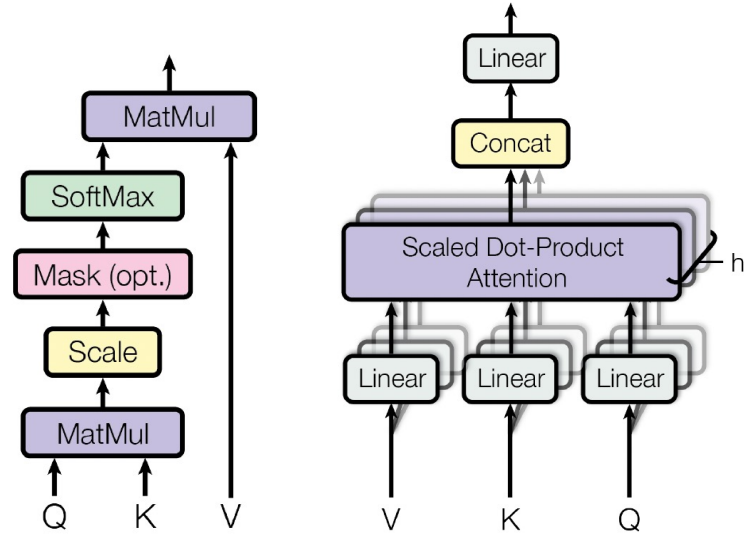



## (2) FNO-Variants

*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

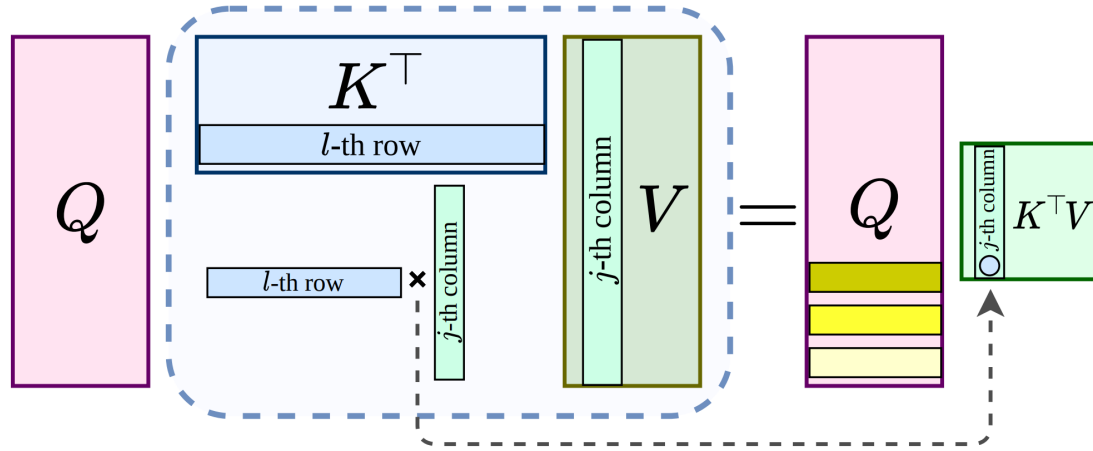


**(1) Geometries as point sequences (2) Attention as Monte Carlo Integral**

*OFormer, Galerkin Transformer, etc*

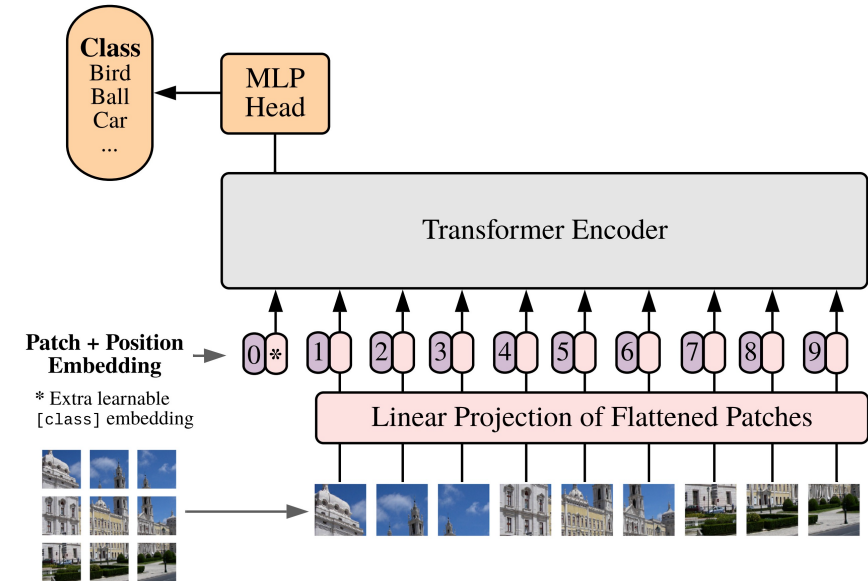
***How to efficiently capture physical correlations underlying discretized meshes is the key to “transform” Transformers into practical PDE solvers***

# 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

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

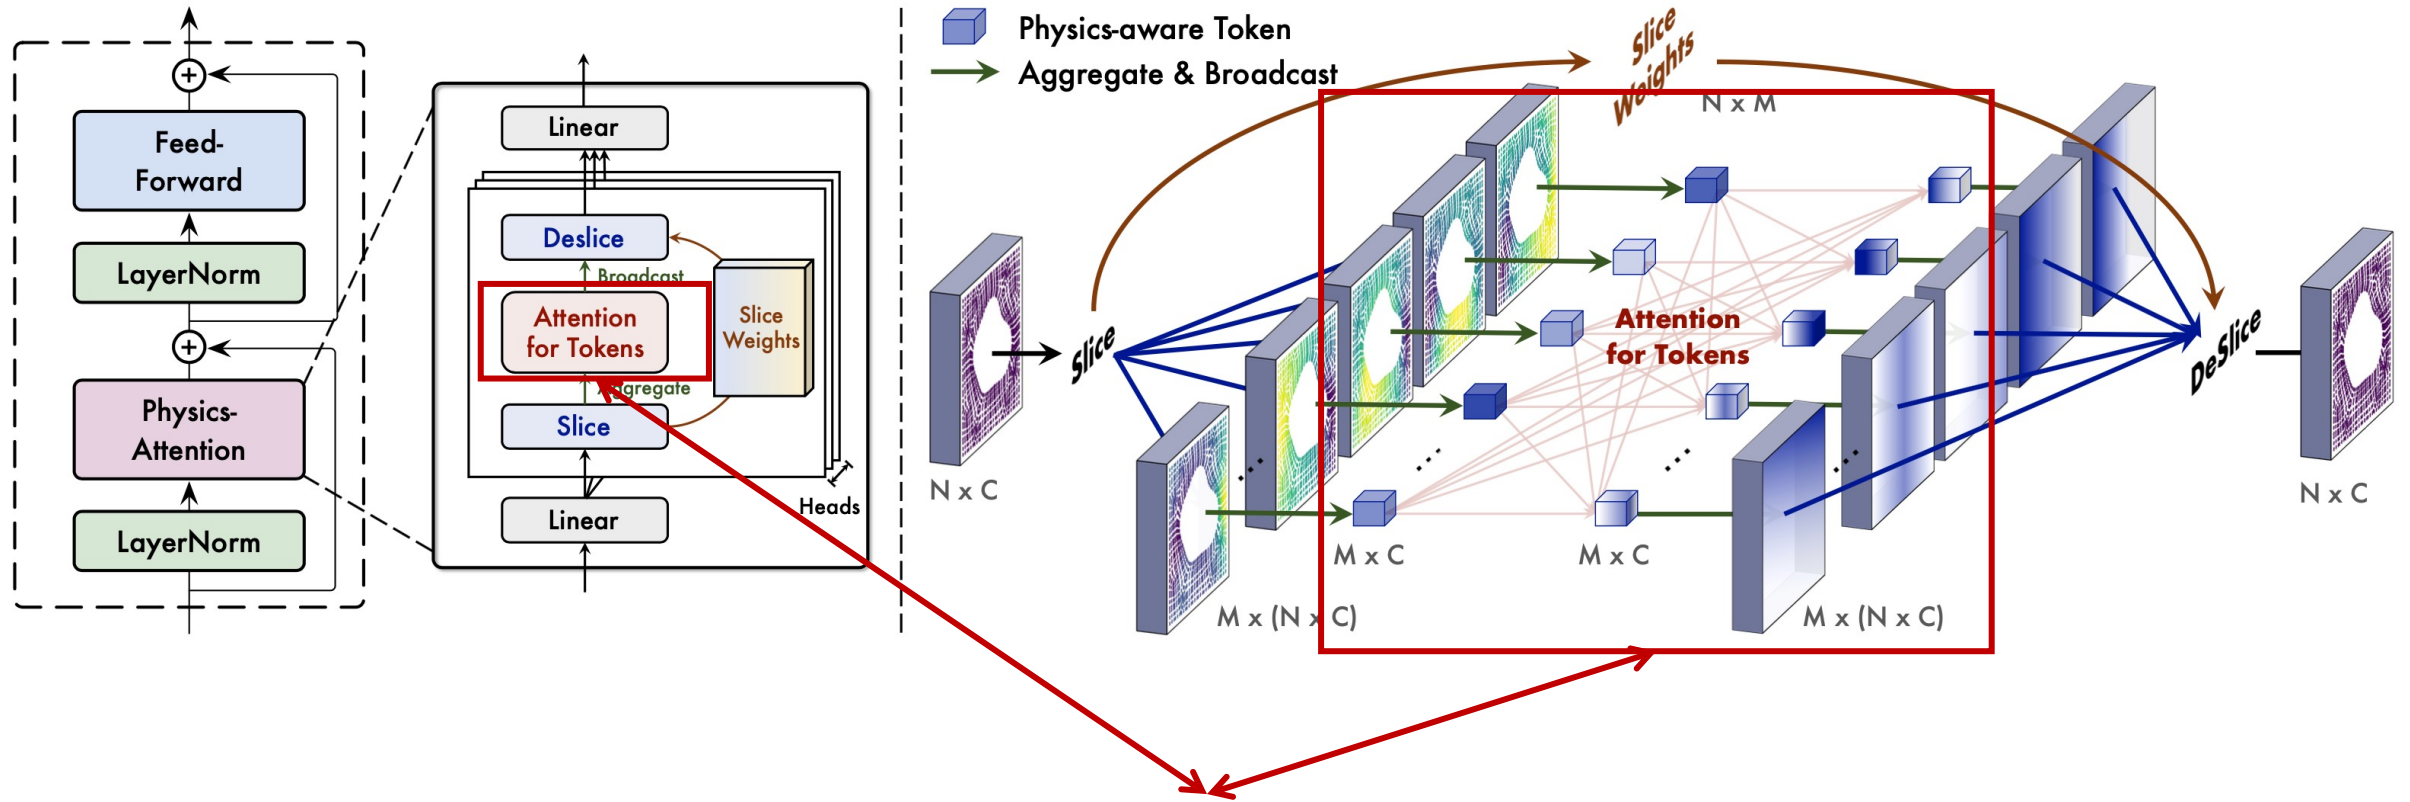
# Mesh $\rightarrow$ physics



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

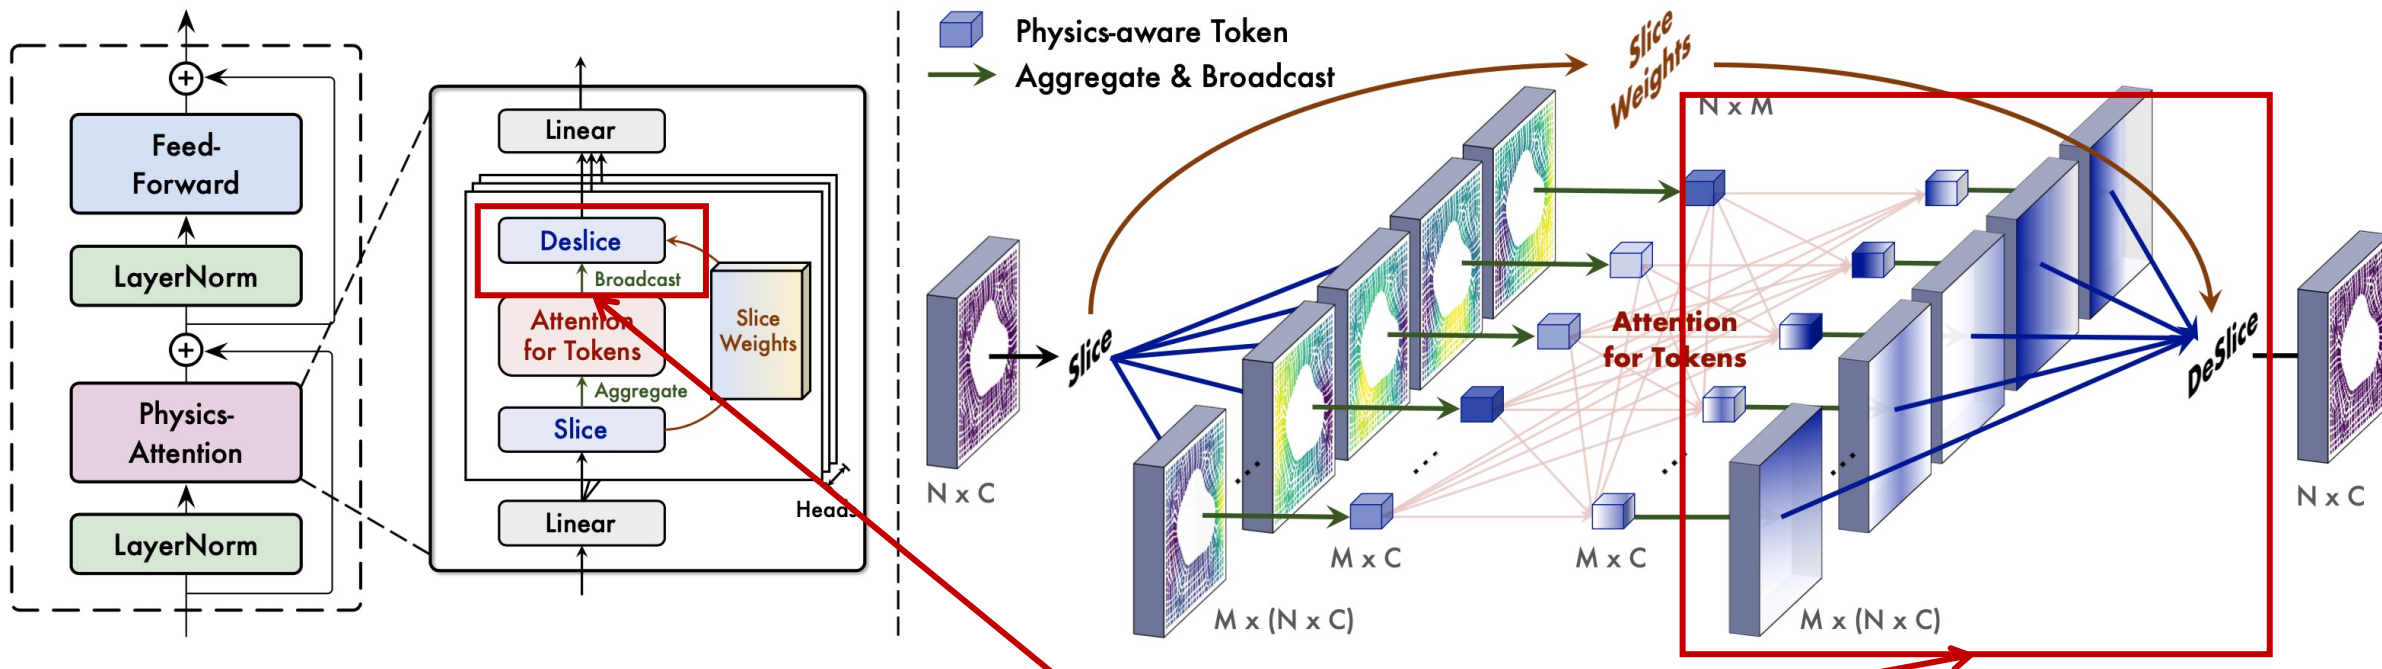


$$\mathbf{q}, \mathbf{k}, \mathbf{v} = \text{Linear}(\underline{\mathbf{z}}), \quad \mathbf{z}' = \text{Softmax} \left( \frac{\mathbf{q}\mathbf{k}^\top}{\sqrt{C}} \right) \mathbf{v}$$

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
<b>TRANSOLVER (OURS)</b>	<b>0.0064</b>	<b>0.0012</b>	<b>0.0053</b>	<b>0.0033</b>	<b>0.0900</b>	<b>0.0057</b>
RELATIVE PROMOTION	25.6%	29.4%	10.2%	29.7%	24.7%	12.3%

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

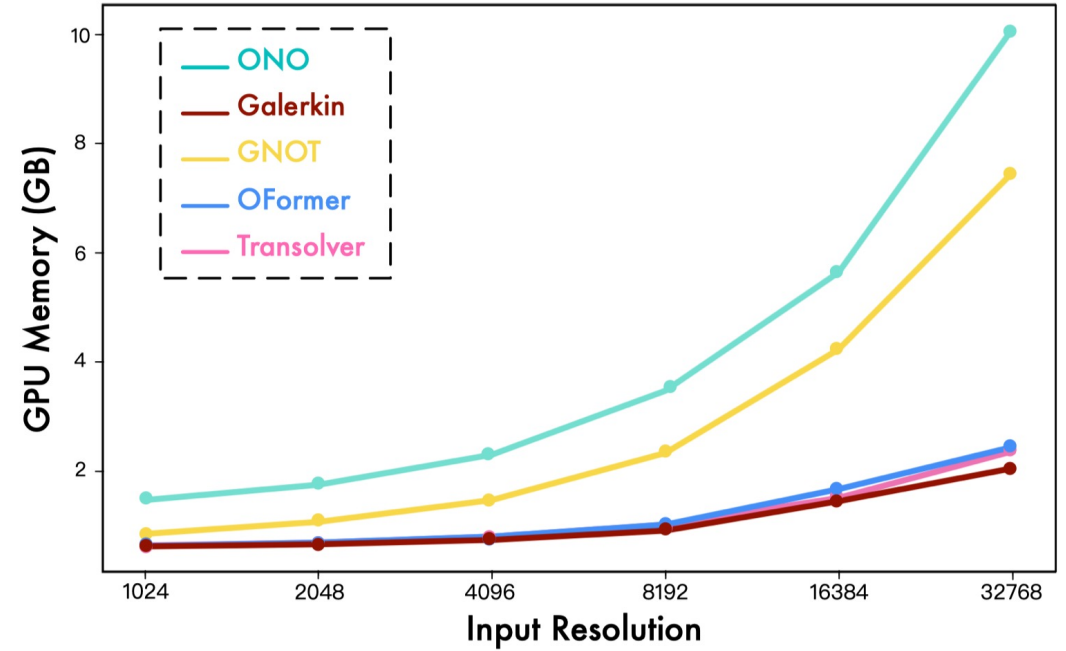

# Practical Design Tasks

MODEL*	SHAPE-NET CAR				AIRFRANS			
	VOLUME ↓	SURF ↓	<u><math>C_D</math> ↓</u>	<u><math>\rho_D</math> ↑</u>	VOLUME ↓	SURF ↓	<u><math>C_L</math> ↓</u>	<u><math>\rho_L</math> ↑</u>
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	<u>0.1476</u>	<u>0.9964</u>
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	<u>0.0049</u>	<u>0.0152</u>	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)	<u>0.0319</u>	<u>0.0779</u>	<u>0.0159</u>	<u>0.9842</u>	/	/	/	/
<b>TRANSOLVER (OURS)</b>	<b>0.0207</b>	<b>0.0745</b>	<b>0.0103</b>	<b>0.9935</b>	<b>0.0037</b>	<b>0.0142</b>	<b>0.1030</b>	<b>0.9978</b>

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

Transolver  
Attention among Tokens



Galerkin Transformer  
Attention among Mesh Points




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

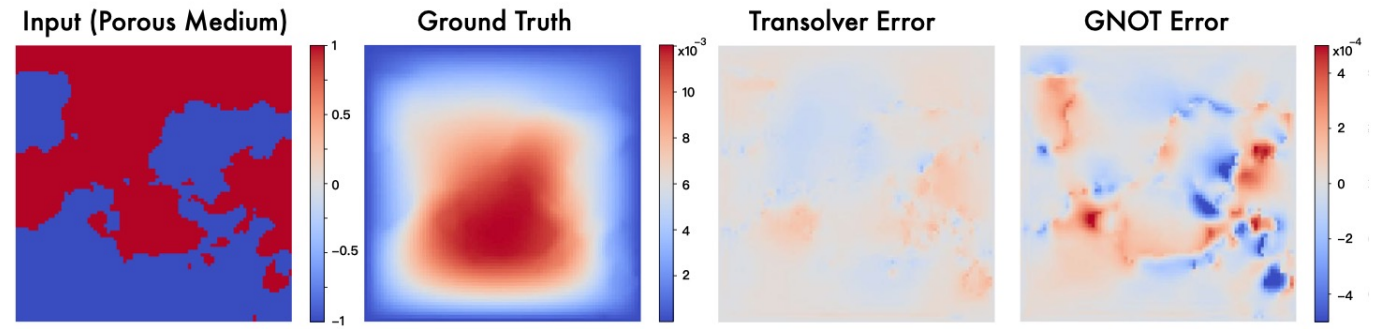
BENCHMARKS	GALERKIN (CAO, 2021)	TRANSOLVER (OURS)
ELASTICITY (972 MESH POINTS)	0.3803	<b>1.7795</b>
DARCY (7,225 MESH POINTS)	0.2739	<b>1.8274</b>

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

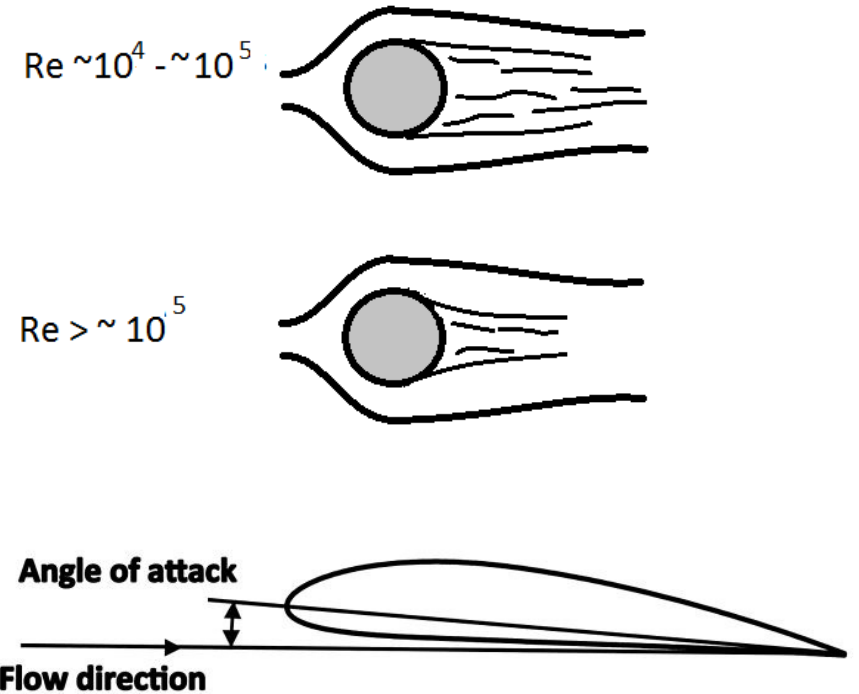


- 1. Resolution:** Consistent performance at varied scales
- 2. Data:** Benefiting from larger training data
- 3. Parameter:** Benefiting from more parameters



# 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>
<b>TRANSOLVER (OURS)</b>	<b>0.2996</b>	<b>0.9896</b>	<b>0.1500</b>	<b>0.9950</b>




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
<b>GNN + TRANSOLVER (OURS)</b>	<b>0.0069</b>
RELATIVE PROMOTION	62.1%

Transolver can also be extended to

**Lagrangian Settings**  
**(Ever-changing geometrics)**

# Open Source

The screenshot shows the GitHub repository page for 'thuml / Transolver'. The repository is public and has 1 branch (main), 0 tags, 7 commits, 3 watchers, 0 forks, and 1 star. The commit history shows updates to files like 'Airfoil-Design-AirFRANS', 'Car-Design-ShapeNetCar', 'PDE-Solving-StandardBenchmark', 'pic', '.gitignore', 'LICENSE', 'Physics\_Attention.py', and 'README.md'. The README section is expanded, showing the title 'Transolver (ICML 2024)' and a description: 'Transolver: A Fast Transformer Solver for PDEs on General Geometries [paper]'. It highlights features such as calculating attention among learned physical states and achieving a 22% error reduction over previous SOTA in six standard benchmarks. The right sidebar contains sections for 'About' (with a link to the arXiv paper), 'Releases', 'Packages', 'Languages' (showing Python at 97.9%, Jupyter Notebook at 1.3%, and Shell at 0.8%), and 'Suggested workflows'.

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



Thank You!

wuhx23@mails.tsinghua.edu.cn



长按关注，获取最新资讯