

# 000 HGsolver: POSITION-ENHANCED PHYSICS AT- 001 TENTION INFORMED HETEROGENEOUS GEOMETRIES 002 NEURAL SOLVER FOR PDES 003

006 **Anonymous authors**

007 Paper under double-blind review

## 011 ABSTRACT

013 Partial differential equations (PDEs) provide a fundamental framework for mod-  
014 eling complex physical phenomena. However, modeling PDEs on heterogeneous  
015 geometries remains a significant challenge for both traditional numerical solvers  
016 and neural operator methods, as sparse observations, multiphysics interactions,  
017 and distinct discretizations often produce heterogeneous geometries between the  
018 observation and output spaces. In this work, we introduce a unified perspective  
019 on physics attention, formulating physical states as projections of observation  
020 embeddings onto learnable functional bases in Hilbert space. Building on this  
021 formulation, we introduce a position-enhanced physics attention mechanism that  
022 incorporates coordinate representations of these bases via rotary position embed-  
023 dings, thereby enabling more effective modeling of heterogeneous interactions.  
024 Leveraging this mechanism, we develop HGsolver, an encoder-decoder frame-  
025 work designed for PDE tasks on heterogeneous domains. Extensive experiments  
026 demonstrate that HGsolver achieves state-of-the-art performance across forward,  
027 inverse, and reconstruction benchmarks under heterogeneous geometries, while  
028 a minimally modified variant, TransolverXP, also delivers competitive results on  
029 standard homogeneous benchmarks. These findings highlight the importance of  
030 effective interactions among physical states in advancing neural PDE solvers and  
031 their potential to address the complexity of the heterogeneous real-world geo-  
032 metries.

## 033 1 INTRODUCTION

034 As a torch, partial differential equations (PDEs) establish a foundational framework for humanity  
035 to gain insight and conquer the physical world through the lens of mathematical models within  
036 the realm of scientific inquiry. Analytically solving most PDEs is infeasible, and conventional nu-  
037 mercial solvers(Zhang et al., 2021) are usually confined to specific classes of PDEs, necessitating  
038 customization for each new formulation. This task-specific nature often leads to limited general-  
039 ity, reduced flexibility, and substantial computational overhead, particularly when handling com-  
040 plex geometries(Umetani & Bickel, 2018) or performing repeated simulations under varying condi-  
041 tions(Umetani & Bickel, 2018). In recent years, deep neural network surrogates, particularly neural  
042 operators (Li et al., 2020; 2023c), have emerged as a computationally efficient paradigm for advanc-  
043 ing the modeling and approximation of complex PDEs. Leveraging their strong nonlinear modeling  
044 capability(Lu et al., 2021), they can learn input–output mappings of PDE-governed tasks from data  
045 and infer solutions much faster than numerical methods.

046 PDEs are typically discretized into large-scale meshes with complex geometries to enable precise  
047 simulations. Several backbones(Hao et al., 2023; Li et al., 2023b) leveraging individual point fea-  
048 tures have been proposed to address these challenges. However, these point-based architectures  
049 encounter significant difficulties in capturing the intricate physical correlations in PDEs, particu-  
050 larly in industrial design involving extremely large and complex geometries, which typically feature  
051 highly coupled multiphysics interactions(Trockman & Kolter, 2022).

052 To this end, Transolver(Wu et al., 2024) introduces a state-based attention mechanism that can oper-  
053 ate under arbitrary general geometries. Nevertheless, it fails to capture the heterogeneity of interac-  
054 tions among state tokens, especially those derived from distinct discretizations. In real-world scenar-

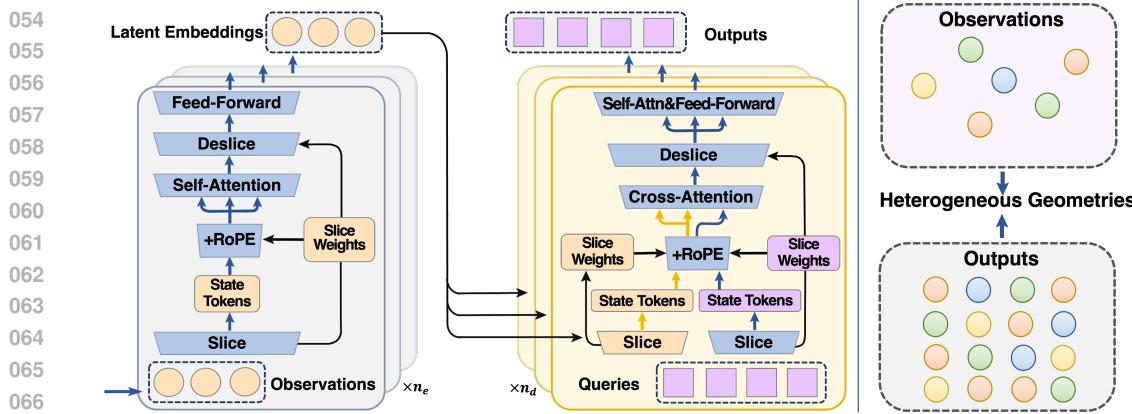


Figure 1: Overview of the **HGsolver** architecture (**left**) and representative illustrations of heterogeneous geometries (**right**).

ios, the available input is often sparse in both spatial and temporal domains, since experimental data are typically collected using a limited set of sensors or measurement devices. Consequently, inputs and outputs frequently come from heterogeneous geometries, while sufficiently complete outputs are crucial for providing accurate guidance in complex applications such as computational fluid dynamics for car and airplane design. Several diffusion-based methods(Gao & Ji, 2019; Huang et al., 2024), which exploit robust generative capabilities, have also been proposed to address this challenge. However, these approaches suffer from substantial drawbacks, including prolonged training time, slow inference, and an excessive number of parameters. These limitations highlight the urgent need for methods that can more effectively handle problems defined on heterogeneous geometries.

In this paper, we present a unified and novel perspective on physics attention, formulated in a setting where physical states are represented as projections of input embeddings onto a set of well-learned functional bases in Hilbert space. Consequently, the distinctions between different physical states are determined by the variations among these learned bases. Building on this foundation, the learned functional bases can be represented by coordinates in Hilbert space, which provide a more precise characterization of their relative positions, particularly when the functional bases are discretized over heterogeneous geometries. To this end, we propose a position-enhanced physics attention mechanism that integrates these coordinate representations into all physical state tokens through rotary position embeddings (Su et al., 2021). As illustrated in Fig. 1, we further introduce **HGsolver**, a framework specifically designed to address PDE tasks on heterogeneous geometries such as super-pixel and restructured domains. With the proposed position-enhanced physics attention, **HGsolver** efficiently captures complex heterogeneous interactions.

To sum up, our main contributions are as follows:

- (i) From the perspective of projecting embeddings onto learnable bases in Hilbert space, we offer a novel insight into physical states and the mechanism of physics attention.
- (ii) Building upon the positional representations of the learnable bases, we introduce a position-enhanced physics attention mechanism that more effectively captures interactions among physical states, particularly for those discretized over heterogeneous geometries.
- (iii) Leveraging cross position-enhanced physics attention, we design **HGsolver** based on an encoder-decoder architecture, enabling it to tackle tasks involving heterogeneous geometries between inputs and outputs, such as sparse prediction and super-pixel domains.
- (iv) **HGsolver** achieves sota performance on benchmarks involving heterogeneous geometries, including forward, inverse, and reconstruction problems of PDEs, while a minimally modified variant **TransolverXP**, also remains competitive on standard benchmarks.

## 2 METHODOLOGY

In this section, we present a novel perspective on physics attention and an effective approach for its enhancement. We begin by reviewing the conventional implementations of physics attention, and then provide a comprehensive description of our methodology and the improvements it introduces.

108  
109

## 2.1 PROBLEM SETUP

110 Our study focuses on two fundamental categories of PDEs: static PDEs and time-dependent (dy-  
111 namic) PDEs. Static PDEs, such as Darcy’s law (Hubbert, 1956) and models that describe the  
112 behavior of solid materials(Dym et al., 1973), are employed to characterize the equilibrium states  
113 of physical systems. In contrast, time-dependent PDEs model the temporal evolution of physical  
114 systems, such as the Navier–Stokes equations governing fluid dynamics(McLean, 2012).

115 In this paper, we investigate PDEs defined on a domain  $\Omega \subset \mathbb{R}^d \times [0, T]$ , with the solution function  
116 space  $\mathcal{H}$  over  $\Omega$ . The objective is to learn an operator  $\mathcal{G}$  that maps the input function space  $\mathcal{A}$  to the  
117 solution space  $\mathcal{H}$ , i.e.,  $\mathcal{G} : \mathcal{A} \rightarrow \mathcal{H}$  and  $\mathcal{G}(a) = u \in \mathcal{H}$  is the solution function over  $\Omega$ .

118 To learn the operator  $\mathcal{G}$ , neural operator model is trained using a dataset  $\mathcal{D} = \{(a_k, u_k)\}_{1 \leq k \leq D}$ ,  
119 where  $u_k = \mathcal{G}(a_k)$ . In practice,  $a_k$  is often discretized on the mesh  $\{\mathbf{x}_i \in \Omega\}_{1 \leq i \leq N}$ . The discretized  
120 representation  $a_k$  is given by  $\{(\mathbf{x}_i, \mathbf{a}_k^i)\}_{1 \leq i \leq N}$ , where  $\mathbf{a}_k^i = a_k(\mathbf{x}_i)$ . In this manner, the input  
121 functions  $a_k$  are represented by  $\{(\mathbf{x}_i, \mathbf{a}_k^i)\}_{1 \leq i \leq N}$ .

122 For the corresponding solution function  $u_k$ , we discretize it on the mesh  $\{\mathbf{y}_i \in \Omega\}_{1 \leq i \leq N'}$ , with  
123 the discretized representation  $u_k$  given by  $\{(\mathbf{y}_i, \mathbf{u}_k^i)\}_{1 \leq i \leq N'}$ , where  $\mathbf{u}_k^i = u_k(\mathbf{y}_i)$ . To model the  
124 operator,  $\mathcal{G}$  is parameterized by a neural network  $\tilde{\mathcal{G}}_w$ , which takes the pair  $(\mathbf{x}_i, \mathbf{a}_k^i)$  as input, and  
125 outputs  $\tilde{\mathcal{G}}_w((\mathbf{x}_i, \mathbf{a}_k^i), \mathbf{y}_i) = \{\tilde{\mathbf{u}}_k^i\}_{1 \leq i \leq N'}$  on mesh  $\{\mathbf{y}_i\}_{1 \leq i \leq N'}$ , which serves as an approximation  
126 to  $u_k$ . The goal is to minimize the mean squared error (MSE) loss between the predicted solution  
127 and the true data, as given by

$$129 \min_{w \in W} \frac{1}{D} \sum_{k=1}^D \frac{1}{N'} \left\| \tilde{\mathcal{G}}_w((\mathbf{x}_i, \mathbf{a}_k^i), \mathbf{y}_i) - \{\mathbf{u}_k^i\}_{1 \leq i \leq N'} \right\|_2^2 \quad (1)$$

132 where  $w$  denotes the set of network parameters and  $W$  represents the space of  $w$ . Detailed structure  
133 of  $\tilde{\mathcal{G}}_w$  will be introduced in following sections.

134  
135 2.2 IMPLEMENTATION OF PHYSICS ATTENTION

136 While attention mechanisms(Vaswani et al., 2017) have demonstrated remarkable effectiveness in  
137 both computer vision (CV)(Khan et al., 2022) and natural language processing (NLP)(Brown et al.,  
138 2020), their application to neural PDE solvers remains challenging. In particular, attention-based  
139 approaches often struggle to faithfully capture the complex physical correlations within the  
140 computational domain  $\Omega$ , which is typically discretized into a set of mesh points  $\{\mathbf{z}_i\}_{i=1}^N$ . These mesh  
141 points, being finite samples of the underlying continuous physical space, motivate the need to learn  
142 the intrinsic physical states. However, the large number of discretized points in the mesh can over-  
143whelm the attention mechanism, making it challenging to identify reliable correlations. To address  
144 this, physics-based attention(Wu et al., 2024), by focusing on physics-sensitive regions, proves to  
145 be particularly effective in capturing the underlying physical information within complex and high-  
146 dimensional meshes.

147 Given a mesh with  $N$  nodes on positions  $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^N$ , to capture the intrinsic physical interactions  
148 under the challenging mesh configurations, Physics-Attention first assigns  $C$ -channel input embed-  
149 dings  $\mathbf{Z} = \{\mathbf{z}_i\}_{i=1}^N \in \mathbb{R}^{N \times C}$  to  $M$  physical states  $\mathbf{S} = \{\mathbf{s}_j\}_{j=1}^M \in \mathbb{R}^{M \times C}$ , based on the learned  
150 slice weights  $\mathbf{W} = \{\mathbf{w}_i\}_{i=1}^N \in \mathbb{R}^{N \times M}$  from the inputs. Each weight vector  $\mathbf{w}_i \in \mathbb{R}^{1 \times M}$  represents  
151 the degree that point  $\mathbf{x}_i$  with embeddings  $\mathbf{z}_i$  belongs to each physical state. Specifically, the physical  
152 states are aggregated from all point representations based on the learned slice weights, which can be  
153 formalized as:

$$154 \text{Physical States: } \mathbf{S} = \{\mathbf{s}_j\}_{j=1}^M = \left\{ \frac{\mathbf{W}_{\cdot, j}^\top \mathbf{Z}}{\mathbf{W}_{\cdot, j}^\top \mathbf{1}_N} \right\}_{j=1}^M, \quad (2)$$

155 The success of Physics-Attention hinges on the assumption of homogeneous physical states. Trans-  
156 solver utilizes the Softmax function to compute the slice weights, thereby ensuring a sharper distri-  
157 bution. However, the phenomenon of degeneration persists as the model depth increases. To mitigate  
158 this issue, Transolver++(Luo et al., 2025) introduces a local adaptive mechanism that employs the  
159 Gumbel-Softmax(Jang et al., 2016) for differentiable sampling from the discrete categorical distri-  
160 bution, and sets the temperature parameter  $\tau_0$  in the Gumbel-Softmax function learnable, as outlined

162 in Eq. (3).  
 163

164 Rep-Slice :  $\mathbf{W} = \{\mathbf{w}_i\}_{i=1}^N = \left\{ \text{Softmax} \left( \frac{\text{Linear}(\mathbf{z}_i) - \log(-\log \epsilon_i)}{\tau_i} \right) \right\}_{i=1}^N$  (3)  
 165  
 166 Ada-Temp:  $\tau = \{\tau_i\}_{i=1}^N = \{\tau_0 + \text{Linear}(\mathbf{z}_i)\}_{i=1}^N$ ,

167 where  $\tau_0$  is the temperature constant and  $\epsilon = \{\epsilon_i\}_{i=1}^N, \epsilon_i \sim \mathcal{U}(0, 1)$ . Here  $\log(-\log \epsilon_i) \sim$   
 168 Gumbel(0, 1), where Gumbel is a type of generalized extreme value distribution. After that, a  
 169 vanilla attention mechanism is applied on  $\mathbf{S} = \{\mathbf{s}_j\}_{j=1}^M$  to capture intricate correlations among  
 170 different states  $\{\mathbf{s}_j\}_{j=1}^M$  as follows.  
 171

172  
 173  $\mathbf{Q}, \mathbf{K}, \mathbf{V} = \text{Linear}(\mathbf{S}), \mathbf{S}' = \text{Softmax} \left( \frac{\mathbf{Q} \mathbf{K}^\top}{\sqrt{C}} \right) \mathbf{V}$ . (4)  
 174  
 175

176 Finally, Physics-Attention applies the deslice operation to project the updated states  $\{\mathbf{s}'_j\}_{j=1}^M$  back  
 177 onto the mesh representation  $\{\mathbf{z}'_i\}_{i=1}^N$  through slice weights as Eq. (5). This operation establishes a  
 178 learnable non-linear mapping between  $\mathbf{Z}$  and  $\mathbf{Z}'$ . Consequently, Physics-Attention can be flexibly  
 179 stacked into more complex architectures like transformer, thereby enhancing its ability to capture  
 180 the underlying physical correlations.  
 181

182  $\mathbf{Z}' = \{\mathbf{z}'_i\}_{i=1}^N = \left\{ \sum_{j=0}^M \mathbf{w}_{i,j} \mathbf{s}'_j \right\}_{i=1}^N$  (5)  
 183  
 184

### 185 2.3 A PROJECTION PERSPECTIVE OF PHYSICS ATTENTION 186

187 As shown in Eq. (2), the  $j$ -th physical state can be expressed as  $\mathbf{s}_j = \mathbf{v}_j \mathbf{Z} / N$ , where  $\mathbf{v}$  is defined  
 188 as  $\mathbf{v}_j := (N \mathbf{W}_{\cdot,j}^\top) / (\mathbf{W}_{\cdot,j}^\top \mathbf{1}_N) \in \mathbb{R}^{1 \times N}$ . Here,  $\mathbf{v}_j := \mathbf{v}(\mathbf{Z}, \theta_j)$  is parameterized by the learnable  
 189 parameters  $\theta_j$  and the input embeddings  $\mathbf{Z} \in \mathbb{R}^{N \times C}$  as described in Eq. (3). The input embedding  
 190  $\mathbf{Z}_{i,\cdot}$  encodes the latent features at position  $\mathbf{x}_i$ . The normalized slice weights  $v_{j,i}$  quantify the fraction  
 191 of features that the state  $\mathbf{s}_j$  receives from position  $\mathbf{x}_i$ . Consequently, for a given position  $\mathbf{x}_k$ , one  
 192 can identify the corresponding input embedding  $\mathbf{Z}_k$  and the slice weights  $v_{j,k}$  defined at position  $\mathbf{x}_i$ .  
 193 Thus, the quantities  $\mathbf{Z}$  and  $\mathbf{v}(\mathbf{Z}, \theta_j)$  can be equivalently represented in a functional form as shown  
 194 in Eq. (6), with the corresponding state  $\mathbf{s}_j$  expressed as in Eq. (7).  
 195

196  $\kappa_{\{\mathbf{Z}, \theta_j\}} : \Omega \rightarrow \mathbb{R}^1, \text{ where } \kappa_{\{\mathbf{Z}, \theta_j\}}(\mathbf{x}_i) = v_{j,i} \quad i = 1, 2, \dots, N$  (6)  
 197  $h : \Omega \rightarrow \mathbb{R}^C, \text{ where } h(\mathbf{x}_i) = \mathbf{Z}_{\{i,\cdot\}}$   $i = 1, 2, \dots, N$

198  
 199  $\mathbf{s}_j = \frac{1}{N} \mathbf{v} \mathbf{Z} = \frac{1}{N} \sum_{i=1}^N \kappa_{\{\mathbf{Z}, \theta_j\}}(\mathbf{x}_i) h(\mathbf{x}_i), \quad j = 1, 2, \dots, M$  (7)  
 200  
 201

202 In the following sections, we simplify the notation by rewriting  $\kappa_{\mathbf{Z}, \theta_j}$  as  $\kappa_j$  for clarity. As demon-  
 203 strated in Eq. (8), the physical state can be interpreted as a numerical approximation of an integral  
 204 of the product of slice weights and input embeddings.  
 205

206  $\mathbf{s}_j = \frac{1}{N} \sum_{i=1}^N \kappa_j(\mathbf{x}_i) h(\mathbf{x}_i) \approx \frac{1}{\|\Omega\|} \int_{\Omega} \kappa_j(\xi) h(\xi) d\xi = \frac{\langle \kappa_j, h \rangle}{\|\Omega\|}, \quad j = 1, 2, \dots, M$  (8)  
 207

208 It worth to note that the physical state  $\mathbf{s}_j$  can be interpreted as the inner product between the input  
 209 function  $h$  and a learnable functional basis  $\kappa_j$  in a Hilbert space. Without loss of generality, we set  
 210  $\|\Omega\| = 1$  for simplicity. Moreover, if we apply an  $L^2$  regularization on  $\mathbf{W}$ , defining as:  
 211

212  $v_j = \{\kappa_j(\mathbf{x}_i)\}_{i=1}^N = (N \mathbf{W}_{\cdot,j}^\top) / (\sqrt{\mathbf{W}_{\cdot,j}^\top \mathbf{W}_{\cdot,j}}), \quad j = 1, 2, \dots, M$  (9)  
 213

214 Then,  $\kappa_j$  becomes a unit vector satisfying  $\langle \kappa_j, \kappa_j \rangle = 1$  in the Hilbert space  $L^2(\Omega)$ . Consequently,  
 215 the physical state  $\mathbf{s}_j$  can be interpreted as the projection of the input function  $h$  onto the learnable  
 216 functional basis  $\kappa_j$  where  $\mathbf{s}_j = \langle \kappa_j, h \rangle / \langle \kappa_j, \kappa_j \rangle$ .  
 217

Moreover, the unnormalized physics attention score between the  $m$ -th and  $n$ -th states, prior to applying the Softmax operation, can be expressed as Eq. (10). Here, ATT denotes the vanilla attention functional score in the original space  $h$  discretized as  $N$  tokens. The attention score is computed by first projecting the state representations onto query and key vectors, where  $\mathbf{W}_q$  and  $\mathbf{W}_k$  are learnable weight matrices associated with the query and key transformations, respectively.

$$\begin{aligned} \text{P-ATT}_{m,n} &= C^{-1/2} \langle \kappa_m, h\mathbf{W}_q \rangle^\top \langle \kappa_n, h\mathbf{W}_k \rangle \\ &= C^{-1/2} \iint \kappa_m(\xi_1) h(\xi_1) \mathbf{W}_q \mathbf{W}_k^\top h(\xi_2)^\top \kappa_n(\xi_2) d\xi_1 d\xi_2 \\ &= C^{-1/2} \iint \kappa_m(\xi_1) \text{ATT}(\xi_1, \xi_2) \kappa_n(\xi_2) d\xi_1 d\xi_2 = C^{-1/2} \langle \kappa_m, \kappa_n \rangle_{\text{ATT}} \end{aligned} \quad (10)$$

As shown in Eq. (10), the attention score in the physics attention mechanism can be interpreted as the inner product of learnable basis functions  $\kappa$ , weighted by the vanilla attention score, which is computed over the original space with  $N$  tokens. To summarize, we now present a projection-based perspective of physical states and a weighted inner-products interpretation on the operation of physics attention. In this interpretation, the slice component is designed to learn a set of essential bases from the input functional  $h$ . By projecting  $h$  onto these learned bases, we obtain distinct physical states. Computing the inner products of the learned bases, modulated by canonical attention scores, effectively captures the interactions among different physical states. However, vanilla Physics Attention primarily attends to the scalar magnitudes of these projections while disregarding the geometric orientations of the bases. As a consequence, distinct bases  $\kappa_1$  and  $\kappa_2$ , corresponding to different physical subspaces of  $h$ , become indistinguishable from the perspective of other states whenever  $\langle \kappa_1, h \rangle = \langle \kappa_2, h \rangle$ , thereby neglecting critical subspace structure.

This limitation reveals that, in its conventional formulation, Physics Attention cannot encode directional information inherent in the learned bases, which may be crucial for distinguishing certain physical states. To overcome this deficiency, it is necessary to endow the state tokens with directional awareness of the learnable bases. Concretely, this is achieved by embedding positional information into the state tokens within the functional space  $L^2(\Omega)$ , thereby allowing the self-attention mechanism to exploit not only projection magnitudes but also geometric orientations. Building on this principle, we propose an enhanced Physics Attention mechanism incorporating positional encoding, which will be detailed in the following sections.

## 2.4 COORDINATE REPRESENTATION OF PHYSICAL STATES

As discussed in Sec. 2.3, Physics Attention is designed to learn a set of robust functional bases  $\{\kappa_1, \dots, \kappa_M\}$  and to compute the projections of the input embeddings  $\mathbf{Z}$  (i.e., the input function  $h$  defined over the domain  $\Omega$ ) onto these bases. A natural and effective way to characterize the relationships among these bases is to represent them in terms of their coordinates with respect to a common set of orthogonal bases, which enables a systematic analysis of their relative configurations.

Similarly, the Physics Cross-Attention mechanism—which performs standard cross-attention between two groups of physical states—can be interpreted from the same projection-based perspective. In this context, it aims to learn two sets of robust functional bases. If the combined set of bases,  $\{\kappa_1^q, \dots, \kappa_{M_q}^q, \kappa_1^k, \dots, \kappa_{M_k}^k\}$ , is defined over the same domain  $\Omega$ , it can likewise be represented in terms of their coordinates with respect to a common orthogonal basis, allowing for a unified and coherent analysis of their relative configurations. In the following parts, we present some methods for constructing such orthogonal bases and for computing the corresponding coordinate representations.

**Functional Principal Component Analysis (FPCA) Method:** A classic approach for computing these coordinates is to employ spectral decomposition within the Hilbert space framework. To derive the coordinate representations using FPCA (Shang, 2014), we first center the set of basis functions  $\{\kappa_1, \dots, \kappa_M\}$  and define the corresponding empirical covariance operator  $\mathcal{R}$  as Eq. (11).

$$\tilde{\kappa}_i = \kappa_i - \frac{1}{M} \sum_{j=1}^M \kappa_j, \quad i = 1, \dots, M, \quad \mathcal{R} = \frac{1}{M} \sum_{j=1}^M \tilde{\kappa}_j \otimes \tilde{\kappa}_j \quad (11)$$

The goal of FPCA is then to identify a function  $\epsilon \in L^2(\Omega)$  with unit norm,  $\|\epsilon\|_{L^2(\Omega)} = 1$ , such that it satisfies the eigenvalue problem as follows:

$$\mathcal{R}\epsilon = \lambda\epsilon, \quad \text{where} \quad \mathcal{R}\epsilon = \frac{1}{M} \sum_{j=1}^M \langle \tilde{\kappa}_j, \epsilon \rangle \tilde{\kappa}_j \quad (12)$$

270 Here,  $\mathcal{R}$  represents the empirical covariance operator acting on the centered functional observations  $\{\tilde{\kappa}_j\}_{j=1}^M$ , and  $\lambda$  denotes the eigenvalue associated with the principal component  $\epsilon$ . Since the operator  $\mathcal{R}$  has rank at most  $M$ , all nonzero eigenfunctions necessarily belong to the subspace  $\text{span}\{\tilde{\kappa}_1, \dots, \tilde{\kappa}_M\}$ . Hence, each eigenfunction  $\epsilon_i$  can be represented as a linear combination of the centered sample functions defined as follows:

$$275 \quad \epsilon_i = \beta_i^\top \tilde{\kappa} = \sum_{j=1}^M \beta_{ij} \tilde{\kappa}_j, \quad i = 1, \dots, M \quad (13)$$

278 where  $\beta_i \in \mathbb{R}^M$  denotes the coefficient vector. Substituting Eq. (12) into the eigenvalue equation  
279  $\mathcal{R}\epsilon = \lambda\epsilon$  and introducing the Gram matrix  $\mathbf{K} \in \mathbb{R}^{M \times M}$  with entries  $\mathbf{K}_{ij} = \langle \tilde{\kappa}_i, \tilde{\kappa}_j \rangle$ , we obtain:

$$281 \quad \lambda_i \beta_i^\top \tilde{\kappa} = \frac{1}{M} \sum_{j=1}^M \langle \tilde{\kappa}_j, \epsilon_i \rangle \tilde{\kappa}_j = \frac{1}{M} \sum_{j=1}^M \sum_{l=1}^M \beta_{il} \langle \tilde{\kappa}_j, \tilde{\kappa}_l \rangle \tilde{\kappa}_j = \frac{1}{M} (\mathbf{K} \beta_i)^\top \tilde{\kappa}, \quad i = 1, \dots, M \quad (14)$$

283 By comparing the leftmost and rightmost terms in Eq. (14), we conclude that  $\frac{1}{M} \mathbf{K} \beta_i = \lambda_i \beta_i$ , which  
284 implies that  $\beta_i$  is an eigenvector of the Gram matrix  $\mathbf{K}$  corresponding to eigenvalue  $M\lambda_i \geq 0$ .  
285 Consequently, the original infinite-dimensional eigenvalue problem in  $L^2(\Omega)$  is reduced to a finite-  
286 dimensional eigenvalue problem for the Gram matrix  $\mathbf{K}$ , which can be readily solved using standard  
287 linear algebraic techniques.

288 Let  $\eta_i$  denote a unit eigenvector of  $\mathbf{K}$  corresponding to the eigenvalue  $M\lambda_i$ . By setting  $\beta_i = a\eta_i$   
289 with  $a > 0$ , the constant  $a$  is determined by the normalization condition  $|\epsilon_i|_H = 1$ . Specifically,  
290

$$291 \quad \|\epsilon_i\|_H^2 = \beta_i^\top \mathbf{K} \beta_i = a^2 \eta_i^\top (M\lambda_i \eta_i) = a^2 M \lambda_i = 1, \quad \Rightarrow \quad a = (M\lambda_i)^{-1/2}. \quad (15)$$

292 Consequently, the coordinates of the original uncentered functional bases  $\{\kappa_1, \dots, \kappa_M\}$  with respect  
293 to the orthogonal eigenbasis  $\{\epsilon_1, \dots, \epsilon_M\}$  are encoded in the matrix  $\Theta \in \mathbb{R}^{M \times M}$ , defined as  
294

$$295 \quad \Theta_{ij} = \langle \kappa_i, \epsilon_j \rangle = \frac{(\tilde{\kappa}_i + \bar{\kappa})^\top \tilde{\kappa}^\top \eta_j}{\sqrt{M\lambda_j}} = \frac{\mathbf{e}_i^\top \mathbf{K} \eta_j}{\sqrt{M\lambda_j}} + C_j = \sqrt{M\lambda_j} \mathbf{e}_i^\top \eta_j + C_j. \quad (16)$$

297 where  $\mathbf{e}_i$  is the  $i$ -th canonical basis vector with 1 in the  $i$ -th position and 0 elsewhere and  $C_j$  denotes  
298 a constant that does not depend on  $i$ . The computational complexity of this procedure is  $O(N + M^3)$ .  
299 For completeness, we also present two alternative schemes for computing the coordinates. However,  
300 as their empirical performance on the benchmarks proved inferior to that of the aforementioned  
301 method, we provide their descriptions in Appendix A.2 for reference.

## 302 2.5 POSITION-ENHANCED PHYSICS ATTENTION WITH POSITIONAL ENCODING

304 The vanilla attention mechanism is inherently position-agnostic when positional information is not  
305 explicitly incorporated into the input features. To overcome this limitation, the Galerkin/Fourier  
306 Transformer (Cao, 2021) augments each attention head by concatenating spatial coordinates with  
307 latent embeddings, and further applies a spectral convolutional decoder (Li et al., 2020) on top of  
308 the attention layers. Here, we adopt Rotary Position Embeddings (RoPE) (Su et al., 2021) to encode  
309 positional information. Originally proposed for modeling relative positions in language models,  
310 RoPE offers a flexible and effective mechanism for integrating coordinate-dependent information  
311 into the attention framework.

312 Following the discussion in Sec. 2.4, we obtain the  $M$ -dimensional coordinate matrix  $\Theta \in \mathbb{R}^{M \times M}$ ,  
313 which represents the  $M$  physical states  $\mathbf{s}_{i=1}^M$ . Since  $M$  is typically set to 32 or higher, directly  
314 embedding them into tokens is often impractical. In addition, not all coordinate components are rel-  
315 evant to the underlying physical interactions among states. To address this, we introduce a learnable  
316 linear transformation that projects the  $M$ -dimensional coordinates into a lower-dimensional latent  
317 space as  $\Theta \leftarrow \mathbf{T}\Theta$ . Importantly, the projection matrix  $\mathbf{T}$  is shared across the entire model.

318 For a given latent coordinate vector  $\theta = \Theta_{i, \cdot}$  and its corresponding  $H$ -dimensional latent embedding  
319  $\mathbf{q} = S_{i, \cdot} \mathbf{W}_q$ , we select the latent dimension  $H$  such that  $H/(2J) = \gamma \in \mathbb{N}^+$ , and accordingly  
320 reshape  $\mathbf{q}$  into a matrix  $\tilde{\mathbf{q}} \in \mathbb{R}^{J \times 2\gamma}$ . Subsequently, we define the RoPE operator  $\phi(\text{Coor,Emb})$  as:

$$321 \quad \phi(\mathbf{q}, \theta) = \text{Concat}(\mathbf{R}(\theta_1)\tilde{\mathbf{q}}_1, \dots, \mathbf{R}(\theta_J)\tilde{\mathbf{q}}_J), \quad \mathbf{R}(\theta_j) = \text{diag}(\mathbf{R}_1(\theta_j), \dots, \mathbf{R}_\gamma(\theta_j))$$

$$323 \quad R_k(\theta_j) = \begin{bmatrix} \cos(\theta_j \nu_k) & -\sin(\theta_j \nu_k) \\ \sin(\theta_j \nu_k) & \cos(\theta_j \nu_k) \end{bmatrix} \text{ where } \nu_k = 10000^{-(k-1)/\gamma}, k = 1, \dots, \gamma \quad (17)$$

324 where `diag` denotes the construction of a block-diagonal matrix by placing the sub-matrices along its  
 325 diagonal, and the frequency coefficients  $\nu_k$  are defined as in (Vaswani et al., 2017; Su et al., 2021).  
 326 The operator  $\Phi(\cdot)$  is then naturally defined by applying the  $\phi(\cdot)$  operator row-wise across the entire  
 327 coordinate matrix.

328 We summarize our proposed **Position-Enhanced Physics Attention** (PPA) as follows. For **Self-  
 329 attention** (Self-PPA) and **Cross-attention** (Cross-PPA), Eq. (4) is adapted into Eq. (18) and  
 330 Eq. (19), respectively, as an extension of the vanilla physics attention(Wu et al., 2024).  
 331

$$332 \quad \mathbf{Q}, \mathbf{K}, \mathbf{V} = \text{Linear}(\mathbf{S}), \quad \Theta = \text{Linear}[\text{Get\_Coor}(\mathbf{W})], \\ 333 \quad \mathbf{S}' = \text{Softmax} \left( \frac{\Phi(\mathbf{Q}, \Theta) \Phi(\mathbf{K}, \Theta)^\top}{\sqrt{C}} \right) \mathbf{V} \quad (18) \\ 334 \\ 335$$

$$336 \quad \mathbf{Q} = \text{Linear}(\mathbf{S}_q), \quad \mathbf{K}, \mathbf{V} = \text{Linear}(\mathbf{S}_k), \quad \Theta_q, \Theta_k = \text{Linear}[\text{Get\_Coor}(\mathbf{W}_q, \mathbf{W}_k)], \\ 337 \quad \mathbf{S}' = \text{Softmax} \left( \frac{\Phi(\mathbf{Q}, \Theta_q) \Phi(\mathbf{K}, \Theta_k)^\top}{\sqrt{C}} \right) \mathbf{V} \quad (19) \\ 338 \\ 339$$

340 Here, the procedure `Get_Coor` is defined as in Sec. 2.4.  
 341

### 342 3 MODEL ARCHITECTURE

343 **Query Encoder** Given a query position  $\mathbf{y}_i$ , the query encoder employs a shared point-wise MLP  
 344 whose first layer implements a random Fourier feature mapping  $\mathcal{Q}(\cdot)$  (Tancik et al., 2020; Rahimi  
 345 & Recht, 2007). The Gaussian random Fourier mapping is defined as:  
 346

$$347 \quad \text{Query} := \mathcal{Q}(\mathbf{Y}) = \text{Concat} [\cos(2\pi \mathbf{Y} \mathbf{B}), \sin(2\pi \mathbf{Y} \mathbf{B})], \quad (20) \\ 348$$

349 where  $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_{n'}]^\top \in \mathbb{R}^{N' \times \dim(\Omega)}$  and  $\mathbf{y}_i \in \mathbb{R}^d$  denotes the Cartesian coordinates of the  
 350  $i$ -th query point. The projection matrix  $\mathbf{B} \in \mathbb{R}^{d \times d'}$  has entries sampled i.i.d. from the Gaussian  
 351 distribution  $\mathcal{N}(0, \sigma^2)$ , where  $\sigma$  is a predefined scale parameter. The concatenation of cosine and  
 352 sine components yields a  $2d'$ -dimensional embedding per point. By mapping input coordinates to a  
 353 trigonometric basis with higher-frequency components,  $\mathcal{Q}(\cdot)$  serves to mitigate the spectral bias of  
 354 coordinate-based neural networks (Tancik et al., 2020; Mildenhall et al., 2020); analogous encodings  
 355 have also been incorporated into physics-informed machine learning (Wang et al., 2021).

356 **Design on Heterogeneous Geometry** For operator learning tasks defined on heterogeneous  
 357 geometries (i.e., when the input and output are discretized differently), we adopt an encoder-decoder  
 358 backbone based on the canonical transformer architecture (Vaswani et al., 2017), as illustrated in  
 359 Fig. 1. The detailed implementation is summarized in Eq. (21). We refer to this modified archi-  
 360 tecture as **HGsolver**. In the formulation below,  $n_e$  and  $n_d$  denote the depths of the encoder and  
 361 decoder, respectively, and FF represents the feedforward neural network.  
 362

$$362 \quad \text{KV}^{(l+1)} = (\text{FF} \circ \text{Self-PPA})(\text{KV}^{(l)}), \quad l = 0, \dots, n_e - 1, \quad \text{KV}^{(0)} = \text{Input}, \\ 363 \\ 364 \quad \text{Out}^{(l+1)} = (\text{FF} \circ \text{Cross-PPA})(\text{Out}^{(l)}, \text{KV}^{(n_e)}), \quad l = 0, \dots, n_d - 1, \quad \text{Out}^{(0)} = \text{Query}. \quad (21) \\ 365$$

366 In Cross\_PPA, we highlight that the query basis  $\kappa_i$  and the key basis  $\kappa_j$  may be discretized on  
 367 distinct meshes,  $\mathbf{X} \in \mathbb{R}^{N \times \dim(\Omega)}$  and  $\mathbf{Y} \in \mathbb{R}^{N' \times \dim(\Omega)}$ , respectively. Under this condition, the  
 368 inner product  $\langle \kappa_i, \kappa_j \rangle$  is computed using a learnable attention matrix  $\mathbf{L} \in \mathbb{R}^{N \times N'}$ , which serves  
 369 as the kernel matrix for the inner product, as described in Eq. (22). Here,  $\mathbf{L}$  is shared across the  
 370 HGsolver, and  $\mathbf{v}$  represents the slice weights defined in Sec. 2.3.  
 371

$$372 \quad \langle \kappa_i, \kappa_j \rangle = (\mathbf{v}_i^{(q)})^\top \mathbf{L} (\mathbf{v}_j^{(k)}), \quad \mathbf{L} = \text{Softmax} \left( \frac{(\mathbf{X} \mathbf{W}_q^{(\text{pos})})(\mathbf{Y} \mathbf{W}_k^{(\text{pos})})^\top}{\dim(\Omega)} \right) \quad (22) \\ 373$$

374 **Design on Homogeneous Geometry** In operator learning tasks performed on homogeneous ge-  
 375ometries (i.e., where the input and output share the same shape and exhibit a one-to-one mapping  
 376 between corresponding positions), we extend the architecture of Transolver (Wu et al., 2024) by  
 377 replacing its standard physics attention mechanism with our position-enhanced physics attention.  
 This modified architecture is referred to as **TransolverXP**.

378 Table 1: Relative  $\ell_2$  errors of HGsolver on inverse (**left**), reconstruction (**middle**), and forward  
 379 (**right**) tasks. *w/o PE* denotes HGsolver with vanilla physics attention. For the Shape-Net car  
 380 benchmark, the  $\ell_2$  errors of the surrounding physical fields are reported.

Sparse-Ratio	Model	N-S	Elasticity	Darcy	Airfoil	Pipe	ShN-Car
10%	HG	0.28	0.08	0.04	0.19	0.16	0.23
	<i>w/o PE</i>	0.31	0.15	0.05	0.19	0.18	0.29
30%	HG	0.24	0.08	0.02	0.17	0.06	0.18
	<i>w/o PE</i>	0.24	0.14	0.02	0.18	0.09	0.22
50%	HG	0.19	0.07	0.01	0.10	0.04	0.16
	<i>w/o PE</i>	0.21	0.12	0.01	0.13	0.06	0.20

## 4 EXPERIMENTS

In this section, we evaluate HGsolver on heterogeneous geometry benchmarks and TransolverXP on homogeneous ones. These benchmarks span different structures and dimensions, and cover forward, inverse, and reconstruction tasks, providing a convincing evaluation of our proposed models.

**Datasets:** Our experiments encompass a broad range of problem settings across both 2D and 3D domains, including point clouds (Elasticity), structured meshes (Plasticity, Airfoil, Pipe), regular grids (Navier–Stokes, Darcy), and unstructured meshes (ShapeNet Car, AirRANS). The Elasticity, Plasticity, Airfoil, Pipe, Navier–Stokes, and Darcy benchmarks were first introduced in FNO Li et al. (2021) and geo-FNO Li et al. (2022), and have since become widely adopted in subsequent studies. Beyond these standard benchmarks, we further consider design-oriented tasks: the ShapeNet Car dataset Umetani & Bickel (2018), which involves predicting surface pressure and surrounding air velocity from vehicle geometries, and the AirRANS dataset Bonnet et al. (2022), which provides high-fidelity Reynolds-Averaged Navier–Stokes simulations of airfoils derived from the National Advisory Committee for Aeronautics. In addition, we include an inverse problem benchmark introduced by LNO (Wang & Wang, 2024), involving the solution of the inverse problem for the Burgers equation. Detailed descriptions and specifications of all the aforementioned benchmarks are provided in Appendix A.1.

**Baselines:** We conduct a comprehensive comparison of HGsolver against more than 20 baseline models. These include typical neural operators such as DeepOnet (Lu et al., 2021), FNO (Li et al., 2021), U-NO (Rahman et al., 2023), and LSM (Wu et al., 2023); Transformer-based PDE solvers such as GNOT (Hao et al., 2023) and FactFormer (Li et al., 2023a); as well as classical geometric deep models including PointNet (Qi et al., 2017), GraphSAGE (Hamilton et al., 2017), and Mesh-GraphNet (Pfaff et al., 2021). Among these, LSM (Wu et al., 2023) and GNOT (Hao et al., 2023) represent the previous sota on standard benchmarks, while GINO Li et al. (2023b) and 3D-GeoCA (Deng et al., 2024) are advanced models designed for large-scale, industrial-level simulation benchmarks. In addition, we also compare with recent physics-attention-based models, including Transolver (Wu et al., 2024), LNO (Wang & Wang, 2024), and Transolver++ (Luo et al., 2025).

**Setup:** We conducted our experiments within an open source framework *Neural Solver Library* (Wu et al., 2024). All experiments were run on 8 Nvidia RTX4090 GPUs with 24GB memory.

### 4.1 STANDARD HETEROGENEOUS GEOMETRIES BENCHMARKS

To assess the impact of incorporating positional encoding into physics attention, we design a comprehensive set of PDE tasks encompassing forward, inverse, and reconstruction processes. These tasks provide a rigorous and convincing evaluation of our approach, with detailed implementations presented in Appendix A.1.2. As shown in Tab. 1, physics attention enhanced with the proposed positional encoding demonstrates superior performance across this diverse set of tasks, particularly when handling complex heterogeneous geometries, such as those in the Shape-Net Car and Elasticity benchmarks. Compared to vanilla physics attention, our enhanced approach more effectively captures the intrinsic physical interactions across a wide range of PDE scenarios.

### 4.2 HETEROGENEOUS GEOMETRIES BENCHMARKS ON INVERSE PROBLEM

To ensure a rigorous and fair comparison, we evaluate HGsolver against state-of-the-art models for PDE tasks on heterogeneous geometries, using the inverse problem benchmark of the Burgers equation shown in Appendix A.1.2. As shown in Tab. 2, the proposed HGsolver consistently demon-

Table 2: Relative MAE of inner reconstructions from *completers* under different observation ratios **(left)**, and of outlier reconstructions from *propagators* using either ground truth(G.T.) or reconstructed inputs from different observation ratios **(right)**.

Task	Completer					Propagator		
	Models	20%	10%	5%	1%	0.5%	G.T.	10%
DeepONet	2.51%	2.59%	2.82%	3.25%	4.82%	7.34%	11.14%	13.87%
GNOT	1.12%	1.39%	1.62%	1.63%	2.56%	5.45%	8.04%	9.91%
LNO	0.60%	0.74%	0.77%	1.18%	<b>2.05%</b>	3.73%	5.69%	7.72%
<b>HGsolver(Ours)</b>	<b>0.52%</b>	<b>0.59%</b>	<b>0.67%</b>	<b>1.16%</b>	2.11%	<b>3.55%</b>	<b>5.61%</b>	<b>7.69%</b>

Table 3: Performance comparison on design-oriented tasks. Besides the relative  $\ell_2$  error of the surrounding (Volume) and surface (Surf) physics fields, the relative  $\ell_2$  errors of the drag coefficient ( $C_D$ ) and lift coefficient ( $C_L$ ) are also recorded, along with their Spearman’s rank correlations  $\rho_D$  and  $\rho_L$ . For clarity, the best results are in **bold** and complete table is provided in Tab 6.

MODEL*	SHAPE-NET CAR				AIRFRANS			
	VOLUME ↓	SURF ↓	$C_D$ ↓	$\rho_D$ ↑	VOLUME ↓	SURF ↓	$C_L$ ↓	$\rho_L$ ↑
...	...	...	...	...	...	...	...	...
3D-GEOCA	0.0319	0.0779	0.0159	0.9842	/	/	/	/
TRANSOLVER	0.0207	0.0745	0.0103	0.9935	0.0037	0.0142	<b>0.1030</b>	<b>0.9978</b>
<b>TRANSOLVERXP</b>	<b>0.0199</b>	<b>0.0622</b>	<b>0.0095</b>	<b>0.9938</b>	<b>0.0031</b>	<b>0.0094</b>	0.1535	0.9963

strates state-of-the-art performance across varying observation ratios, both in its roles as *propagators* and *completers*, thereby highlighting the decoupling capabilities of HGsolver.

### 4.3 STANDARD HOMOGENEOUS GEOMETRIES BENCHMARKS

We further evaluate **TransolverXP** on eight homogeneous geometry benchmarks. As shown in Tab. 3 and Tab. 5, **TransolverXP** consistently demonstrates competitive performance across these benchmarks, with particularly notable results on the unstructured and complex cases (ShapeNet-Car and AirfRANS), highlighting its effectiveness in handling challenging geometries.

### 4.4 EFFICIENCY ANALYSIS

To evaluate the feasibility and scalability of our proposed method, a critical analysis involves examining the execution time of the eigenvalue problem discussed in Sec. 2.4. Theoretically, the computational complexity of this algorithm is  $O(N + M^3)$ . In practice, for example in the Airfoil case, the total inference time is 0.55 seconds, of which 0.38(70%) seconds is spent on solving the eigenvalue problem. This corresponds to  $M = 64$  and  $N = 11,271$ . A detailed analysis of execution time of the eigenvalue problem is shown in Appendix A.4. As a side note, since our model introduces only negligible additional parameters, its GPU memory consumption remains virtually the same as that of Transolver(Wu et al., 2024) when the network depth and width are fixed.

## 5 CONCLUSION

**Insights** In this work, we introduced HGsolver, a position-enhanced physics attention framework for modeling PDEs on heterogeneous geometries. Our experimental results highlight the pivotal importance of explicitly capturing interactions among physical states to faithfully represent intrinsic dynamics in heterogeneous domains. Moreover, we present a unified projection-based perspective on physics attention, which not only provides a rigorous theoretical foundation for understanding physical states and the principles underlying their interactions but also offers principled guidance for the development of future neural PDE solvers in physical state spaces. We envisage that the proposed framework, together with its underlying theoretical foundation, will significantly advance the practical resolution of complex PDE problems defined on heterogeneous geometries.

**Limitations and Future Work** A primary limitation of HGsolver is its computational efficiency, as the eigenvalue computation dominates inference time. This bottleneck may be exacerbated in large-scale or high-resolution scenarios. Future work will investigate potential approximations and solver optimizations to alleviate the runtime bottleneck, while aiming to maintain accuracy and generalizability across diverse PDE tasks and heterogeneous geometries.

486 REFERENCES  
487

- 488 Åke Björck. Numerics of gram-schmidt orthogonalization. *Linear Algebra and Its Applications*,  
489 197:297–316, 1994.
- 490 Florent Bonnet, Jocelyn Ahmed Mazari, Paola Cinnella, and patrick gallinari. AirfRANS: High  
491 fidelity computational fluid dynamics dataset for approximating reynolds-averaged navier–stokes  
492 solutions. In *NeurIPS Datasets and Benchmarks Track*, 2022.
- 493 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,  
494 Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are  
495 few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- 496 Shuhao Cao. Choose a transformer: Fourier or galerkin. *CoRR*, abs/2105.14995, 2021. URL  
497 <https://arxiv.org/abs/2105.14995>.
- 498 Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li,  
499 Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d  
500 model repository. *arXiv preprint arXiv:1512.03012*, 2015.
- 501 Jingyang Deng, Xingjian Li, Haoyi Xiong, Xiaoguang Hu, and Jinwen Ma. Geometry-guided con-  
502 ditional adaption for surrogate models of large-scale 3d PDEs on arbitrary geometries. In *IJCAI*,  
503 2024.
- 504 Clive L Dym, Irving Herman Shames, et al. *Solid mechanics*. Springer, 1973.
- 505 Hongyang Gao and Shuiwang Ji. Graph u-nets. In *ICML*, 2019.
- 506 Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs.  
507 *NeurIPS*, 2017.
- 508 Zhongkai Hao, Chengyang Ying, Zhengyi Wang, Hang Su, Yinpeng Dong, Songming Liu,  
509 Ze Cheng, Jun Zhu, and Jian Song. Gnot: A general neural operator transformer for operator  
510 learning. *ICML*, 2023.
- 511 Jiahe Huang, Guandao Yang, Zichen Wang, and Jeong Joon Park. Diffusionpde: Generative pde-  
512 solving under partial observation. *arXiv preprint arXiv:2406.17763*, 2024.
- 513 M King Hubbert. Darcy’s law and the field equations of the flow of underground fluids. *Transactions  
514 of the AIME*, 207(01):222–239, 1956.
- 515 Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv  
516 preprint arXiv:1611.01144*, 2016.
- 517 Salman Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and  
518 Mubarak Shah. Transformers in vision: A survey. *ACM computing surveys (CSUR)*, 54(10s):  
519 1–41, 2022.
- 520 Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- 521 Zijie Li, Dule Shu, and Amir Barati Farimani. Scalable transformer for pde surrogate modeling.  
522 *NeurIPS*, 2023a.
- 523 Zong-Yi Li, Daniel Z. Huang, Burigede Liu, and Anima Anandkumar. Fourier neural operator with  
524 learned deformations for pdes on general geometries. *arXiv preprint arXiv:2207.05209*, 2022.
- 525 Zongyi Li, Nikola B. Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya,  
526 Andrew M. Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial dif-  
527 ferential equations. *CoRR*, abs/2010.08895, 2020. URL <https://arxiv.org/abs/2010.08895>.
- 528 Zongyi Li, Nikola Borislavov Kovachki, Kamyar Azizzadenesheli, Burigede liu, Kaushik Bhat-  
529 tacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial  
530 differential equations. In *ICLR*, 2021.

- 540 Zongyi Li, Nikola Borislavov Kovachki, Chris Choy, Boyi Li, Jean Kossaifi, Shourya Prakash Otta,  
 541 Mohammad Amin Nabian, Maximilian Stadler, Christian Hundt, Kamyar Azizzadenesheli, and  
 542 Anima Anandkumar. Geometry-informed neural operator for large-scale 3d PDEs. In *NeurIPS*,  
 543 2023b.
- 544 Zongyi Li, Nikola Borislavov Kovachki, Chris Choy, Boyi Li, Jean Kossaifi, Shourya Prakash Otta,  
 545 Mohammad Amin Nabian, Maximilian Stadler, Christian Hundt, Kamyar Azizzadenesheli, et al.  
 546 Geometry-informed neural operator for large-scale 3d pdes. *arXiv preprint arXiv:2309.00583*,  
 547 2023c.
- 548 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019.
- 549 Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning  
 550 nonlinear operators via deeponet based on the universal approximation theorem of operators.  
 551 *Nature machine intelligence*, 3(3):218–229, 2021.
- 552 Huakun Luo, Haixu Wu, Hang Zhou, Lanxiang Xing, Yichen Di, Jianmin Wang, and Mingsheng  
 553 Long. Transolver++: An accurate neural solver for pdes on million-scale geometries. *arXiv  
 554 preprint arXiv:2502.02414*, 2025.
- 555 Barnes W McCormick. *Aerodynamics, aeronautics, and flight mechanics*. John Wiley & Sons, 1994.
- 556 Doug McLean. Continuum fluid mechanics and the navier-stokes equations. *Understanding Aero-  
 557 dynamics: Arguing from the Real Physics*, 2012.
- 558 Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and  
 559 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis, 2020. URL  
 560 <https://arxiv.org/abs/2003.08934>.
- 561 Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context en-  
 562 coders: Feature learning by inpainting. In *Proceedings of the IEEE/CVF Conference on Computer  
 563 Vision and Pattern Recognition (CVPR)*, 2016.
- 564 Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter Battaglia. Learning mesh-based  
 565 simulation with graph networks. In *ICLR*, 2021.
- 566 Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets  
 567 for 3d classification and segmentation. In *CVPR*, 2017.
- 568 Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In J. Platt,  
 569 D. Koller, Y. Singer, and S. Roweis (eds.), *Advances in Neural Information Processing Systems*,  
 570 volume 20. Curran Associates, Inc., 2007. URL <https://proceedings.neurips.cc/paper/2007/file/013a006f03dbc5392effeb8f18fda755-Paper.pdf>.
- 571 Md Ashiqur Rahman, Zachary E Ross, and Kamyar Azizzadenesheli. U-no: U-shaped neural oper-  
 572 ators. *TMLR*, 2023.
- 573 Mark Sabini and Gili Rusak. Painting outside the box: Image outpainting with gans. *arXiv preprint  
 574 arXiv:1808.08483*, 2018.
- 575 Han Lin Shang. A survey of functional principal component analysis. *ASyA Advances in Statistical  
 576 Analysis*, 98(2):121–142, Apr 2014. ISSN 1863-818X. doi: 10.1007/s10182-013-0213-1. URL  
 577 <https://doi.org/10.1007/s10182-013-0213-1>.
- 578 Charles Spearman. The proof and measurement of association between two things. 1961.
- 579 Jianlin Su, Yu Lu, Shengfeng Pan, Bo Wen, and Yunfeng Liu. Roformer: Enhanced transformer  
 580 with rotary position embedding. *CoRR*, abs/2104.09864, 2021. URL <https://arxiv.org/abs/2104.09864>.
- 581 Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan,  
 582 Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let  
 583 networks learn high frequency functions in low dimensional domains, 2020. URL <https://arxiv.org/abs/2006.10739>.

- 594    Asher Trockman and J Zico Kolter. Patches are all you need? *arXiv preprint arXiv:2201.09792*,  
 595    2022.
- 596
- 597    Nobuyuki Umetani and Bernd Bickel. Learning three-dimensional flow for interactive aerodynamic  
 598    design. *ACM Transactions on Graphics (TOG)*, 2018.
- 599
- 600    Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,  
 601    Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural informa-*  

602    *tion processing systems*, 30, 2017.

603

604    Sifan Wang, Hanwen Wang, and Paris Perdikaris. On the eigenvector bias of fourier feature net-  
 605    works: From regression to solving multi-scale pdes with physics-informed neural networks. *Com-*  

606    *puter Methods in Applied Mechanics and Engineering*, 384:113938, 2021. ISSN 0045-7825. doi:  
 607    <https://doi.org/10.1016/j.cma.2021.113938>. URL <https://www.sciencedirect.com/science/article/pii/S0045782521002759>.

608

609    Tian Wang and Chuang Wang. Latent neural operator for solving forward and inverse pde problems.  
 610    *Advances in Neural Information Processing Systems*, 37:33085–33107, 2024.

611

612    Haixu Wu, Tengge Hu, Huakun Luo, Jianmin Wang, and Mingsheng Long. Solving high-  
 613    dimensional pdes with latent spectral models. In *ICML*, 2023.

614

615    Haixu Wu, Huakun Luo, Haowen Wang, Jianmin Wang, and Mingsheng Long. Transolver: A fast  
 616    transformer solver for pdes on general geometries. *arXiv preprint arXiv:2402.02366*, 2024.

617

618    Zipeng Xiao, Zhongkai Hao, Bokai Lin, Zhijie Deng, and Hang Su. Improved operator learning by  
 619    orthogonal attention. In *ICML*, 2024.

620

621    Zongxin Yang, Jian Dong, Ping Liu, Yi Yang, and Shuicheng Yan. Very long natural scenery im-  
 622    age prediction by outpainting. In *Proceedings of the IEEE/CVF International Conference on*  

623    *Computer Vision (ICCV)*, 2019.

624

625    Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image  
 626    inpainting with contextual attention. In *Proceedings of the IEEE/CVF Conference on Computer*  

627    *Vision and Pattern Recognition (CVPR)*, 2018.

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648 **A APPENDIX**  
649650 **A.1 DATASETS SETTINGS**  
651652 In this section, we provide the details of our experiments, including **benchmarks**, **metrics**, and  
653 **implementations**.  
654655 **A.1.1 DESCRIPTIONS OF STANDARD BENCHMARKS**  
656657 We extensively evaluate our model in eight benchmarks, whose information is summarized in Table  
658 4. Note that these benchmarks involve the following three types of PDEs:  
659

- 660
- **Solid material** (Dym et al., 1973): Elasticity and Plasticity.
  
661   - **Navier-Stokes equations for fluid** (McLean, 2012): Airfoil, Pipe, Navier-Stokes, Shape-  
662 Net Car and AirfRANS.
  
663   - **Darcy’s law** (Hubbert, 1956): Darcy.
- 
- 664

665 Here are the details of each benchmark.  
666667 **Elasticity** This benchmark is to estimate the inner stress of the elasticity material based on the  
668 material structure, which is discretized in 972 points (Li et al., 2022). For each case, the input is a  
669 tensor in the shape of  $972 \times 2$ , which contains the 2D position of each discretized point. The output  
670 is the stress of each point, thus in the shape of  $972 \times 1$ . As for the experiment, 1000 samples with  
671 different structures are generated for training and another 200 samples are used for test.  
672673 **Plasticity** This benchmark is to predict the future deformation of the plasticity material under the  
674 impact from above by an arbitrary-shaped die (Li et al., 2022). For each case, the input is the shape  
675 of the die, which is discretized into the structured mesh and recorded as a tensor with shape  $101 \times 31$ .  
676 The output is the deformation of each mesh point in the future 20 time steps, that is a tensor in the  
677 shape of  $20 \times 101 \times 31 \times 4$ , which contains the deformation in four directions. Experimentally, 900  
678 samples with different die shapes are used for model training and 80 new samples are for test.  
679680 **Airfoil** This task is to estimate the Mach number based on the airfoil shape, where the input shape  
681 is discretized into structured mesh with shape  $221 \times 51$  and the output is the Mach number for each  
682 mesh point (Li et al., 2022). Here, all the shapes are deformed from the NACA-0012 case provided  
683 by the National Advisory Committee for Aeronautics. 1000 samples in different airfoil designs are  
684 used for training and the other 200 samples are for testing.  
685686 **Pipe** This benchmark is to estimate the horizontal fluid velocity based on the pipe structure (Li  
687 et al., 2022). Each case discretizes the pipe into structured mesh with size  $129 \times 129$ . Thus, for each  
688 case, the input tensor is in the shape of  $129 \times 129 \times 2$ , which contains the position of each discretized  
689 mesh point. The output is the velocity value for each point, thus in the shape of  $129 \times 129 \times 1$ . 1000  
690 samples with different pipe shapes are used for model training and 200 new samples are for test,  
691 which are generated by controlling the centerline of the pipe.  
692693 Table 4: Summary of experiment benchmarks, where the first six datasets are from FNO (Li et al.,  
694 2021) and geo-FNO (Li et al., 2022), Shape-Net Car is from (Umetani & Bickel, 2018) and pre-  
695 processed by (Deng et al., 2024), and AirfRANS is from (Bonnet et al., 2022). #Mesh records the size  
696 of discretized meshes. #Dataset is organized as the number of samples in training and test sets.  
697698 

GEOMETRY	BENCHMARKS	#DIM	#MESH	#INPUT	#OUTPUT	#DATASET
POINT CLOUD	ELASTICITY	2D	972	STRUCTURE	INNER STRESS	(1000, 200)
STRUCTURED MESH	PLASTICITY AIRFOIL PIPE	2D+TIME 2D 2D	3,131 11,271 16,641	EXTERNAL FORCE STRUCTURE STRUCTURE	MESH DISPLACEMENT MACH NUMBER FLUID VELOCITY	(900, 80) (1000, 200) (1000, 200)
REGULAR GRID	NAVIER-STOKES DARCY	2D+TIME 2D	4,096 7,225	PAST VELOCITY POROUS MEDIUM	FUTURE VELOCITY FLUID PRESSURE	(1000, 200) (1000, 200)
UNSTRUCTURED MESH	SHAPE-NET CAR AIRFRANS	3D 2D	32,186 32,000	STRUCTURE STRUCTURE	VELOCITY & PRESSURE VELOCITY & PRESSURE	(789, 100) (800, 200)

**702 Navier-Stokes** This benchmark is to model the incompressible and viscous flow on a unit torus,  
 703 where the fluid density is constant and viscosity is set as  $10^{-5}$  (Li et al., 2021). The fluid field is  
 704 discretized into  $64 \times 64$  regular grid. The task is to predict the fluid in the next 10 steps based on  
 705 the observations in the past 10 steps. 1000 fluids with different initial conditions are generated for  
 706 training, and 200 new samples are used for test.

**707 Darcy** This benchmark is to model the flow of fluid through a porous medium (Li et al., 2021).  
 708 Experimentally, the process is discretized into a  $421 \times 421$  regular grid. Then we downsample the  
 709 data into  $85 \times 85$  resolution for main experiments. The input of the model is the structure of the  
 710 porous medium and the output is the fluid pressure for each grid. 1000 samples are used for training  
 711 and 200 samples are generated for test, where different cases contain different medium structures.

**712 Shape-Net Car** This benchmark focuses on the drag coefficient estimation for the driving car,  
 713 which is essential for car design. Overall, 889 samples with different car shapes are generated to  
 714 simulate the 72 km/h speed driving situation (Umetani & Bickel, 2018), where the car shapes are  
 715 from the “car” category of ShapeNet (Chang et al., 2015). Concretely, they discretize the whole  
 716 space into unstructured mesh with 32,186 mesh points and record the air around the car and the  
 717 pressure over the surface. Here we follow the experiment setting in 3D-GeoCA (Deng et al., 2024),  
 718 which takes 789 samples for training and the other 100 samples for testing. The input mesh of each  
 719 sample is also preprocessed into the combination of mesh point position, signed distance function  
 720 and normal vector. The model is trained to predict the velocity and pressure value for each point.  
 721 Afterward, we can calculate the drag coefficient based on these estimated physics fields.

**722 AirFRANS** This dataset contains the high-fidelity simulation data for Reynolds-Averaged  
 723 Navier-Stokes equations (Bonnet et al., 2022), which is also used to assist airfoil design. Dif-  
 724 ferent from Airfoil (Li et al., 2022), this benchmark involves more diverse airfoil shapes under finer  
 725 discretized meshes. Specifically, it adopts airfoils in the 4 and 5 digits series of the National Ad-  
 726 visory Committee for Aeronautics, which have been widely used historically. Each case is discretized  
 727 into 32,000 mesh points. By changing the airfoil shape, Reynolds number, and angle of attack, Air-  
 728 FRANS provides 1000 samples, where 800 samples are used for training and 200 for the test set. Air  
 729 velocity, pressure and viscosity are recorded for surrounding space and pressure is recorded for the  
 730 surface. Note that both drag and lift coefficients can be calculated based on these physics quantities.  
 731 However, as their original paper stated, air velocity is hard to estimate for airplanes, making all the  
 732 deep models fail in drag coefficient estimation (Bonnet et al., 2022). Thus, in the main text, we  
 733 focus on the lift coefficient estimation and the pressure quantity on the volume and surface, which  
 734 is essential to the take-off and landing stages of airplanes.

### 735 A.1.2 IMPLEMENTATIONS OF PDE TASKS

**736 Navier-Stokes (inverse)** This inverse task aims to infer the initial fluid state from partially ob-  
 737 served future states. The input consists of velocity and pressure fields at a subset of grid points that  
 738 are masked over the next 10 time steps on a  $64 \times 64$  regular grid. The output corresponds to the  
 739 complete initial fluid state at full spatial resolution. This inverse task evaluates the model’s ability  
 740 to reconstruct physically consistent initial conditions from limited, masked temporal observations.

**741 Elasticity (reconstruction)** This reconstruction task aims to predict the complete stress field from  
 742 partially observed stress values. The input consists of stress values at unmasked points of the  $972 \times 1$   
 743 discretized grid, and the output corresponds to the stress values at the masked locations. This task  
 744 evaluates the model’s ability to infer the full stress distribution from partially observed measure-  
 745 ments.

**746 Darcy (reconstruction)** This reconstruction task involves estimating the fluid pressure at masked  
 747 grid points given partially observed pressures. The input consists of pressures at unmasked points on  
 748 the downsampled  $421 \times 421$  grid, and the output corresponds to the pressure values at the masked  
 749 locations. This task assesses the model’s capability to reconstruct the complete pressure field in  
 750 heterogeneous porous media from limited observations.

**751 Pipe (forward)** This forward task aims to predict the horizontal fluid velocity from the pipe geom-  
 752 etry. The input consists of the spatial coordinates of all  $129 \times 129$  mesh points, with a subset masked  
 753 according to a predefined ratio, and the output corresponds to the velocity at each point, represented

756 as a tensor of shape  $129 \times 129 \times 1$ . This task evaluates the model’s ability to map partially observed  
 757 geometric configurations to the resulting fluid dynamics.  
 758

759 **Airfoil (forward)** This forward task aims to predict the Mach number distribution based on the  
 760 discretized airfoil shape. The input consists of coordinates of a subset of the  $221 \times 51$  structured  
 761 mesh, masked according to a predefined ratio, and the output corresponds to the Mach number at  
 762 each mesh point. This task examines the model’s capability to infer aerodynamic properties from  
 763 partially observed geometric features.  
 764

765 **Shape-Net Car (forward)** This forward task aims to predict the velocity and pressure fields  
 766 around a car from its 3D geometry. The input consists of a preprocessed representation of 32,186  
 767 mesh points, including positions, signed distance functions, and normal vectors, with a subset  
 768 masked according to a predefined ratio. The output corresponds to the velocity and pressure values  
 769 for all points. This task evaluates the model’s effectiveness in estimating complex fluid interactions  
 770 from partially observed 3D vehicle geometries under simulated driving conditions.  
 771

### 772 A.1.3 INVERSE PROBLEM BENCHMARKS

773 Here, we will introduce an inverse benchmark from LNO(Wang & Wang, 2024) as below.

774 To demonstrate the flexibility of the model, an inverse problem is designed. Given a partially ob-  
 775 served solution  $u(x)$ , the objective is to recover the complete solution  $u(x)$  over a larger domain.  
 776 Specifically, the test is conducted on the 1D Burgers’ equation:  
 777

$$778 \quad \frac{\partial}{\partial t} u(x, t) = 0.01 \frac{\partial^2}{\partial x^2} u(x, t) - u(x, t) \frac{\partial}{\partial x} u(x, t), \quad x \in [0, 1], t \in [0, 1] \\ 779 \quad u(x, 0) \sim \text{GaussianProcess}(0, \exp(-\frac{2}{pl^2} \sin^2(\pi ||x - x'||))), \quad u(0, t) = u(1, t)$$

780 The ground-truth data is generated on a  $128 \times 128$  grid with periodic boundary conditions. Initial  
 781 conditions are sampled from a Gaussian process with periodic length  $p = 1$  and scaling factor  $l = 1$ .  
 782

783 The goal of this inverse problem is to reconstruct the complete solution  $u(x)$  across the entire spa-  
 784 tiotemporal domain  $(x, t) \in [0, 1] \times [0, 1]$ , based on sparsely random-sampled or fixed-sampled  
 785 observations in the subdomain  $(x, t) \in [0, 1] \times [0.25, 0.75]$ .  
 786

787 Instead of using a naive approach that directly learns the mapping from partially observed samples  
 788 in the subdomain to the complete solution in the whole domain, a two-stage strategy is proposed,  
 789 inspired by inpainting(Pathak et al., 2016; Yu et al., 2018) and outpainting(Yang et al., 2019; Sabini  
 790 & Rusak, 2018). First, the model is trained as a completer to interpolate the sparsely sampled  
 791 points in the subdomain  $[0, 1] \times [0.25, 0.75]$  to predict all densely and regularly sampled points  
 792 in the same subdomain. Then, the model is trained as a propagator to extrapolate the results of  
 793 the completer from the subdomain to the whole domain  $[0, 1] \times [0, 1]$ . Since the observation and  
 794 prediction samples are located in different positions, only models with decoupling properties can be  
 795 used as the completer and propagator.  
 796

797 The performance of the model is compared with that of DeepONet(Lu et al., 2021), GNOT(Hao  
 798 et al., 2023) and LNO(Wang & Wang, 2024) in both stages.  
 799

### 800 A.1.4 METRICS

801 Since our experiment consists of standard benchmarks and practical design tasks, we also include  
 802 several design-oriented metrics in addition to the relative L2 for physics fields.  
 803

804 **Relative L2 for physics fields** Given the physics field  $\mathbf{u}$  and the model predicted field  $\hat{\mathbf{u}}$ , the  
 805 relative L2 of model prediction can be calculated as follows:  
 806

$$807 \quad \text{Relative L2} = \frac{\|\mathbf{u} - \hat{\mathbf{u}}\|}{\|\mathbf{u}\|}. \quad (23)$$

810    **Mean Absolute Error (MAE)** Given the physics field  $\mathbf{u}$  and the model predicted field  $\hat{\mathbf{u}}$ , the mean  
 811    absolute error of the model prediction can be calculated as follows:  
 812

$$813 \quad \text{MAE} = \frac{1}{N} \sum_{i=1}^N |u_i - \hat{u}_i|, \quad (24)$$

$$814$$

$$815$$

816    **Relative L2 for drag and lift coefficients** For Shape-Net Car and AirfRANS, we also calculated  
 817    the drag and lift coefficients based on the estimated physics fields. For unit density fluid, the coeffi-  
 818    cient (drag or lift) is defined as follows:  
 819

$$820 \quad C = \frac{2}{v^2 A} \left( \int_{\partial\Omega} p(\xi) (\hat{n}(\xi) \cdot \hat{i}(\xi)) d\xi + \int_{\partial\Omega} \tau(\xi) \cdot \hat{i}(\xi) d\xi \right), \quad (25)$$

$$821$$

$$822$$

823    where  $v$  is the speed of the inlet flow,  $A$  is the reference area,  $\partial\Omega$  is the object surface,  $p$  denotes  
 824    the pressure function,  $\hat{n}$  means the outward unit normal vector of the surface,  $\hat{i}$  is the direction  
 825    of the inlet flow and  $\tau$  denotes wall shear stress on the surface.  $\tau$  can be calculated from the air  
 826    velocity near the surface (McCormick, 1994), which is usually much smaller than the pressure item.  
 827    Specifically, for the drag coefficient of Shape-Net Car,  $\hat{i}$  is set as  $(-1, 0, 0)$  and  $A$  is the area of  
 828    the smallest rectangle enclosing the front of cars. As for the lift coefficient of AirfRANS,  $\hat{i}$  is set  
 829    as  $(0, 0, -1)$ . The relative L2 is defined between the ground truth coefficient and the coefficient  
 830    calculated from the predicted velocity and pressure field.

831    **Spearman’s rank correlations for drag and lift coefficients** Given  $K$  samples in the test set with  
 832    the ground truth coefficients  $C = \{C^1, \dots, C^K\}$  (drag or lift) and the model predicted coefficients  
 833     $\hat{C} = \{\hat{C}^1, \dots, \hat{C}^K\}$ , the Spearman correlation coefficient is defined as the Pearson correlation  
 834    coefficient between the rank variables, that is:  
 835

$$836 \quad \rho = \frac{\text{cov} (R(C)R(\hat{C}))}{\sigma_{R(C)}\sigma_{R(\hat{C})}}, \quad (26)$$

$$837$$

$$838$$

839    where  $R$  is the ranking function,  $\text{cov}$  denotes the covariance and  $\sigma$  represents the standard deviation  
 840    of the rank variables. Thus, this metric is highly correlated to the model guide for design opti-  
 841    mization. A higher correlation value indicates that it is easier to find the best design following the  
 842    model-predicted coefficients (Spearman, 1961).  
 843

## 844    A.2 ADDITIONAL METHOD FOR COORDINATES REPRESENTATION

$$845$$

846    **Gram-Schmidt Method:** One straightforward method for computing these coordinates is to ap-  
 847    ply the Gram-Schmidt orthogonalization procedure (Björck, 1994) to the original bases, as illus-  
 848    trated in Eq. (28). Specifically, given a set of bases  $\{\kappa_1, \dots, \kappa_M\}$ , the associated orthogonal bases  
 849     $\{\epsilon_1, \dots, \epsilon_M\}$  can be obtained as follows.

$$850 \quad \epsilon_1 = \frac{\kappa_1}{\|\kappa_1\|}, \quad \|\kappa_1\| = \sqrt{\langle \kappa_1, \kappa_1 \rangle},$$

$$851$$

$$852 \quad \tilde{\epsilon}_i = \kappa_i - \sum_{j=1}^{i-1} \langle \kappa_i, \epsilon_j \rangle \epsilon_j, \quad \epsilon_i = \frac{\tilde{\epsilon}_i}{\sqrt{\langle \tilde{\epsilon}_i, \tilde{\epsilon}_i \rangle}}, \quad i = 2, \dots, M$$

$$853$$

$$854$$

$$855$$

856    Consequently, the coordinates of the original bases  $\kappa$  with respect to the orthogonal bases  $\epsilon$  can be  
 857    represented by the matrix  $\Theta \in \mathbb{R}^{M \times M}$ , where  
 858

$$859 \quad \Theta_{i,j} = \langle \kappa_i, \epsilon_j \rangle \approx \frac{1}{N} \sum_{k=1}^N \kappa_i(x_k) \epsilon_j(x_k), \quad i, j = 1, \dots, M. \quad (28)$$

$$860$$

$$861$$

862    Each row  $\Theta_i$  can thus be interpreted as the “position” of the state  $s_i$  in the space defined by the  
 863    orthogonal bases. The computational complexity of this procedure is  $O(NM^2)$ , which may become  
 864    prohibitively expensive when  $M$  is large.

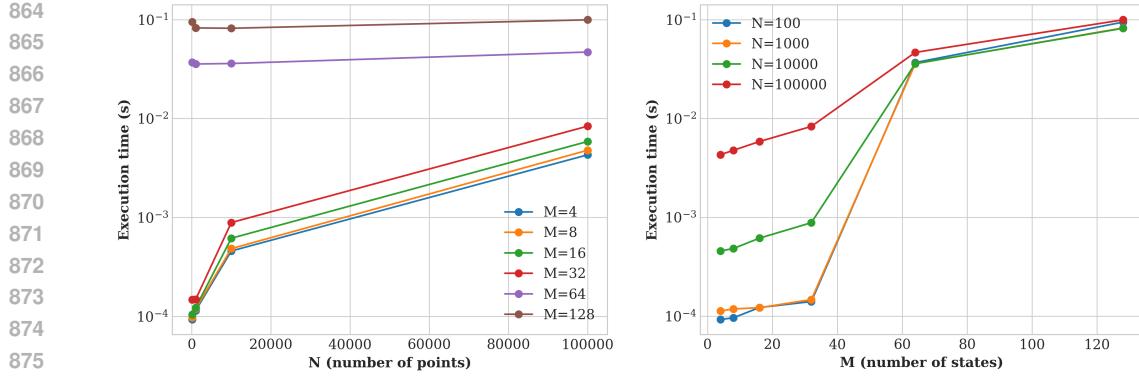


Figure 2: Times Efficiency Analysis on Number of States and Points

**Static Bases Method:** Although the aforementioned methods provide a principled approach for computing the coordinates, empirical strategies are also worth considering, as they typically require fewer computational resources. In this approach, a set of bases  $\{\epsilon_1, \dots, \epsilon_E\}$  is selected empirically to compute the coordinates  $\Theta$  of  $\{\kappa_1, \dots, \kappa_M\}$  as follows:

$$\Theta_{ij} = \langle \kappa_i, \epsilon_j \rangle \in \mathbb{R}^{M \times E}. \quad (29)$$

Subsequently, the columns of  $\Theta$  with the top- $k$  variances are retained and taken as the coordinates of the state tokens. For instance, one may employ a series of Fourier basis functions, as their coefficients can be efficiently computed via the Fast Fourier Transform (FFT). This approach enables the evaluation of the coordinates of  $N$  candidate bases with respect to  $M$  functions at an overall computational cost of  $O(MN \log N)$ .

### A.3 ADDITIONAL EXPERIMENTAL RESULTS

The results of **TransolverXP** on standard PDE benchmark is shown in Tab. 5.

### A.4 TIMES EFFICACY ANALYSIS

The computational complexity of the Gram eigenvalue algorithm is  $O(N + M^3)$ , where  $N$  denotes the number of samples and  $M$  represents the feature dimension. The empirical results show that when  $M$  is relatively small, the execution time grows markedly with increasing  $N$ , suggesting that the linear term in the complexity dominates. As  $M$  becomes larger, however, the cubic term  $M^3$  emerges as the principal factor, leading to a significant increase in runtime while the dependence on  $N$  becomes negligible. This observation is consistent with the theoretical analysis and underscores the scalability limitation of the algorithm with respect to high-dimensional features.

### A.5 IMPLEMENTATION SETTINGS

Table 7 summarizes the training and model configurations of TransolverXP and HGsolver. Training settings, including learning rate, optimizer, batch size, and number of epochs, are adopted directly from previous works (Wu et al., 2024; Bonnet et al., 2022; Hao et al., 2023; Deng et al., 2024) without additional tuning. Loss functions are defined separately for volume ( $\mathcal{L}_v$ ) and surface fields ( $\mathcal{L}_s$ ), with Darcy additionally incorporating a spatial gradient regularization term ( $\mathcal{L}_g$ ) following ONO (Xiao et al., 2024). For HGsolver, we employ a 4+4 encoder-decoder structure, which is applied consistently across all relevant benchmarks.

### A.6 COMPUTER SYSTEM INFORMATION

Detailed information of computer system we conduct our experiments are shown in the *sysinfo.txt* of supplementary materials.

918  
919  
920  
921  
922 Table 5: Relative  $\ell_2$  error on standard benchmarks are presented as mean of 4 runs.“/” means that  
923 the baseline cannot apply to this benchmark. For clarity, the **first**, **second**, and **third** best are high-  
924 lighted. All metrics are derived from Transolve++ (Luo et al., 2025), with models maintaining an  
925 equivalent number of parameters.

MODEL	POINT CLOUD		STRUCTURED MESH			REGULAR GRID	
	ELASTICITY	PLASTICITY	AIRFOIL	PIPE	NAVIER-STOKES	DARCY	
FNO	/	/	/	/	0.1556	0.0108	
WMT	0.0359	0.0076	0.0075	0.0077	0.1541	0.0082	
U-FNO	0.0239	0.0039	0.0269	0.0056	0.2231	0.0183	
GEO-FNO	0.0229	0.0074	0.0138	0.0067	0.1556	0.0108	
U-NO	0.0258	0.0034	0.0078	0.0100	0.1713	0.0113	
F-FNO	0.0263	0.0047	0.0078	0.0070	0.2322	0.0077	
LSM	0.0218	0.0025	0.0059	0.0050	0.1535	0.0065	
GALERKIN	0.0240	0.0120	0.0118	0.0098	0.1401	0.0084	
HT-NET	/	0.0333	0.0065	0.0059	0.1847	0.0079	
OFORMER	0.0183	0.0017	0.0183	0.0168	0.1705	0.0124	
GNOT	0.0086	0.0336	0.0076	0.0047	0.1380	0.0105	
FACTFORMER	/	0.0312	0.0071	0.0060	0.1214	0.0109	
ONO	0.0118	0.0048	0.0061	0.0052	0.1195	0.0076	
TRANSOLVER	<b>0.0064</b>	<b>0.0012</b>	<b>0.0053</b>	0.0033	0.0900	<b>0.0057</b>	
LNO	0.0069	0.0029	<b>0.0053</b>	0.0031	<b>0.0830</b>	<b>0.0063</b>	
TRANSOLVER++	<b>0.0052</b>	<b>0.0011</b>	<b>0.0048</b>	<b>0.0027</b>	<b>0.0719</b>	<b>0.0049</b>	
TRANSOLVERXP	<b>0.0050</b>	<b>0.0015</b>	0.0061	<b>0.0030</b>	<b>0.0842</b>	0.0069	

944  
945  
946  
947  
948  
949  
950 Table 6: Performance comparison on design-oriented tasks is conducted. In addition to the relative  
951  $\ell_2$  error of the surrounding (Volume) and surface (Surf) physics fields, the relative  $\ell_2$  errors of  
952 the drag coefficient ( $C_D$ ) and lift coefficient ( $C_L$ ) are also recorded, along with their correspond-  
953 ing Spearman’s rank correlations  $\rho_D$  and  $\rho_L$ . A Spearman’s correlation value close to 1 indicates  
954 superior performance. For clarity, the best results are in **bold**.  
955

MODEL*	SHAPE-NET CAR				AIRFRANS			
	VOLUME $\downarrow$	SURF $\downarrow$	$C_D \downarrow$	$\rho_D \uparrow$	VOLUME $\downarrow$	SURF $\downarrow$	$C_L \downarrow$	$\rho_L \uparrow$
SIMPLE MLP	0.0512	0.1304	0.0307	0.9496	0.0081	0.0200	0.2108	0.9932
GRAPH SAGE	0.0461	0.1050	0.0270	0.9695	0.0087	0.0184	0.1476	0.9964
POINTNET	0.0494	0.1104	0.0298	0.9583	0.0253	0.0996	0.1973	0.9919
GRAPH U-NET	0.0471	0.1102	0.0226	0.9725	0.0076	0.0144	0.1677	0.9949
MESHGRAPHNET	0.0354	0.0781	0.0168	0.9840	0.0214	0.0387	0.2252	0.9945
GNO	0.0383	0.0815	0.0172	0.9834	0.0269	0.0405	0.2016	0.9938
GALERKIN	0.0339	0.0878	0.0179	0.9764	0.0074	0.0159	0.2336	0.9951
GNOT	0.0329	0.0798	0.0178	0.9833	0.0049	0.0152	0.1992	0.9942
GINO	0.0386	0.0810	0.0184	0.9826	0.0297	0.0482	0.1821	0.9958
3D-GEOCA	0.0319	0.0779	0.0159	0.9842	/	/	/	/
TRANSOLVER	0.0207	0.0745	0.0103	0.9935	0.0037	0.0142	<b>0.1030</b>	<b>0.9978</b>
TRANSOLVERXP	<b>0.0199</b>	<b>0.0622</b>	<b>0.0095</b>	<b>0.9938</b>	<b>0.0031</b>	<b>0.0094</b>	0.1535	0.9963

970  
971

972  
 973  
 974  
 975  
 976  
 977  
 978  
 979  
 980  
 981  
 982  
 983  
 984  
 985  
 986  
 987  
 988  
 989  
 990  
 991

992 Table 7: Training and model configurations of TransolverXP and HGsolver. Training configurations  
 993 are directly from previous works without extra tuning (Wu et al., 2024; Bonnet et al., 2022; Hao  
 994 et al., 2023; Deng et al., 2024). Here  $\mathcal{L}_v$  and  $\mathcal{L}_s$  represent the loss on volume and surface fields  
 995 respectively. As for Darcy, we adopt an additional spatial gradient regularization term  $\mathcal{L}_g$  following  
 996 ONO (Xiao et al., 2024).

997 **Here, we emphasize we apply a 4+4 encoder-decoder structure in HGsolver.**

998  
 999  
 1000  
 1001  
 1002  
 1003  
 1004  
 1005  
 1006  
 1007  
 1008  
 1009  
 1010  
 1011  
 1012  
 1013  
 1014  
 1015  
 1016  
 1017  
 1018  
 1019  
 1020  
 1021  
 1022  
 1023  
 1024  
 1025

BENCHMARKS	TRAINING CONFIGURATION (SHARED IN ALL BASELINES)					MODEL CONFIGURATION			
	LOSS	EPOCHS	INITIAL LR	OPTIMIZER	BATCH SIZE	LAYERS $L$	HEADS	CHANNELS $C$	SLICES $M$
ELASTICITY					1			128	64
PLASTICITY					8			128	64
AIRFOIL	RELATIVE	500	$10^{-3}$	ADAMW (2019)	4			128	64
PIPE	L2				4	8	8	128	64
NAVIER-STOKES					2			256	32
DARCY	$\mathcal{L}_{rL2} + 0.1\mathcal{L}_g$				4			128	64
SHAPE-NET CAR	$\mathcal{L}_v + 0.5\mathcal{L}_s$	200	$10^{-3}$	ADAM (2015)	1			256	32
AIRFRANS	$\mathcal{L}_v + \mathcal{L}_s$	400			1	8	8	256	32