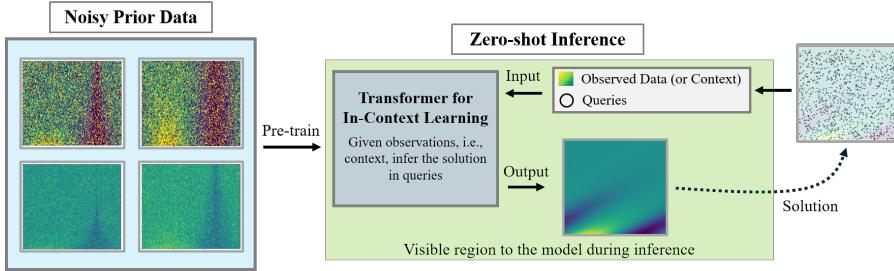


# 000 001 002 003 004 005 PDE-PFN: PRIOR-DATA FITTED NEURAL PDE 006 SOLVER 007 008 009

010 **Anonymous authors**  
011 Paper under double-blind review  
012  
013  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024  
025  
026

## ABSTRACT

027 Despite recent progress in scientific machine learning (SciML), existing approaches  
028 remain impractical, as they often require explicit governing equations, impose rigid  
029 input structures, and lack generalizability across PDEs. Motivated by the success  
030 of large language models (LLMs) with broad generalizability and robustness to  
031 noisy or unreliable pre-training data, we seek to bring similar capabilities to PDE  
032 solvers. In addition, inspired by the Bayesian inference mechanisms of prior-data  
033 fitted networks (PFNs), we propose PDE-PFN, a prior-data fitted neural solver  
034 that directly approximates the posterior predictive distribution (PPD) of PDE  
035 solutions via in-context Bayesian inference. PDE-PFN builds on a PFN architecture  
036 with self- and cross-attention mechanisms of Transformer and is pre-trained on  
037 low-cost approximate solutions generated by physics-informed neural networks,  
038 serving as diverse but not necessarily exact priors. Through experiments on a  
039 range of two-dimensional PDEs, we demonstrate that PDE-PFN achieves strong  
040 generalization across heterogeneous equations, robustness under noisy priors, and  
041 zero-shot inference capability. Our approach not only outperforms task-specific  
042 baselines but also establishes a flexible and robust paradigm for advancing SciML.  
043  
044



045 Figure 1: **The overall workflow of PDE-PFN.** Our model performs in-context learning (ICL) based  
046 on given observations (i.e., context) to infer solutions. Even when trained with an approximated PINN  
047 prior, our method obtains clean solutions due to the implicit Bayesian inference capability of ICL.  
048

## 1 INTRODUCTION

049 Partial differential equations (PDEs) have long been at the core of scientific and engineering research,  
050 as they provide a mathematical description of the fundamental laws governing diverse natural and  
051 engineered systems. Numerical methods such as finite element or finite difference schemes have  
052 traditionally been employed (Quarteroni & Valli, 2008). However, these methods are not only  
053 computationally expensive in high-dimensional or multi-scale settings but also typically require  
054 complete observations of the system. With recent advances in machine learning, scientific machine  
055 learning (SciML) (Raissi et al., 2019a; Willard et al., 2022; Subramanian et al., 2023; Kim et al.,  
056 2023; Choi et al., 2024) has emerged as a promising paradigm for addressing these challenges.

057 Within the trend of SciML, we can see two main approaches: physics-informed neural network  
058 (PINN)-based models (Raissi et al., 2019b; Shukla et al., 2020; Yang et al., 2021; Meng et al., 2020;  
059 Yuan et al., 2022) and neural operator-based models (Li et al., 2020; Lu et al., 2021; Tran et al., 2021;  
060 Guibas et al., 2021). PINN-based models embed physical laws into the learning process by enforcing  
061 PDE residuals as constraints during training, enabling them to approximate solutions without requiring  
062 large labeled datasets. However, a critical limitation of this approach is that it presupposes explicit  
063

knowledge of the underlying PDEs, which is often impractical in many real-world scenarios where governing equations are only partially known or entirely unavailable.

Neural operator-based models aim to learn mappings between infinite-dimensional function spaces, enabling fast inference once trained, and they have achieved notable success on benchmark PDEs. However, their applicability remains limited. First, these models are confined to the specific dataset and task they were trained on, showing poor generalization when applied to new PDEs or distributions. Second, they suffer from input rigidity: they require data to lie on fixed grids or pre-specified coordinates, which prevents flexible inference when observation points vary. These limitations collectively restrict the generalizability and flexibility of neural operator methods, underscoring the need for a more versatile framework that can transfer across diverse PDEs and input conditions.

In recent years, large language models (LLMs) have revolutionized natural language processing by introducing highly flexible and scalable architectures (Brown et al., 2020; Kaplan et al., 2020; Touvron et al., 2023; Frieder et al., 2023; Chowdhery et al., 2023). Motivated by the success of LLMs in enabling broad task generalization, scientific foundation models (SFMs) (McCabe et al., 2023; Yang et al., 2023; Yang & Osher, 2024; Hang et al., 2024) have recently been proposed. These models are typically pre-trained on a wide range of PDEs and subsequently fine-tuned for diverse downstream tasks, extending their applicability beyond the narrow scope of conventional neural operators.

Although SFMs represent a meaningful advance by enabling generalization across different PDEs, they continue to inherit other fundamental drawbacks from their neural operator backbones. In particular, they lack task-level generalization beyond the operator tasks seen during training and impose input rigidity, often constraining data to uniform grids or fixed spatial locations. To address these limitations, we draw inspiration from prior-fitted networks (PFNs) (Müller et al., 2022), which approximate the posterior predictive distribution (PPD) through Bayesian inference conditioned on prior information. Building on this idea, we propose **PDE-PFN**, a new SciML method that adopts a PFN architecture grounded in in-context learning. By directly approximating the PPD of PDE solutions, our method achieves generalizability in terms of the task and PDE type while removing structural constraints on input. In addition, its Bayesian inference enables robustness even when trained with noisy priors, which significantly reduces the data collection costs in SciML.

**Task generalizability on diverse PDEs with enhanced performance** We first evaluated the proposed method on the convection–diffusion–reaction (CDR) equations, a standard benchmark for parameterized PDEs. Our model exhibited strong task generalizability, successfully handling diverse tasks across both parameter and spatiotemporal domains without requiring task-specific fine-tuning. In addition, although pre-trained solely on CDR equations, the model demonstrated PDE generalization by accurately predicting solutions for different PDE families with only minimal fine-tuning. Across these experiments, our method consistently matched or outperformed task-specific baselines and further surpassed existing SFMs in terms of PDE generalizability.

**Flexibility on input shape without prior physical knowledge** Our method advances SciML by removing two key constraints that limit existing approaches. First, it predicts solutions directly from observed quantities, such as velocity and pressure, without requiring access to governing equations. This feature is particularly important in real-world applications where the underlying equations are incomplete or entirely unknown (Chien et al., 2012; Rouf et al., 2021; Beck & Kurz, 2021; Nicolaou et al., 2023; Lee & Cant, 2024). Excluding explicit equations from the input is thus an intentional design choice that broadens applicability across diverse domains. Second, in complex systems such as semiconductor manufacturing, observational data are often collected from sparse or irregularly distributed sensors rather than uniform grids (Myers & Schultz, 2000; Quirk & Serda, 2001; Chien et al., 2012). In such settings, the ability to handle inputs of arbitrary shape is essential. Our method is free from structural constraints on input representation and allows target locations to be flexibly specified as queries. Together, these properties make our approach both practical and versatile, effectively overcoming the limitations of prior SciML models. A detailed comparison with existing baselines is provided in Table 7 of Appendix D.

**Robustness to noisy prior** For LLMs, one of the most challenging steps is collecting prior data, which typically involves crawling and cleaning sentences from the Internet. However, this process is far from perfect due to two key issues: (i) the Internet, as a data source, is inherently unreliable; and (ii) cleaning such vast amounts of data requires significant manual effort. Consequently, LLMs are often trained on incomplete or imperfect prior data. This realistic yet critical issue has been largely overlooked in the literature on SFMs, despite their similarities to LLMs.

108 Notably, numerical solvers are often expensive, time-consuming, and specialized for particular classes  
 109 of PDEs (e.g., the finite-difference time-domain (FDTD) method for Maxwell's equations). In this  
 110 work, we are the first to explore the potential of pre-training a model with PINN-based low-cost,  
 111 noisy, and approximated data, demonstrating that our method can still achieve strong predictive  
 112 performance under such challenging conditions.

113 **Zero-shot inference** Our method supports zero-shot inference for predicting PDE solutions. In  
 114 contrast, ICON-LM (Yang et al., 2025) requires few-shot “demos”<sup>1</sup> for an unknown target operator  
 115 before making predictions. Collecting such demos introduces an inherent delay, since inference  
 116 cannot proceed until these examples are available (see Figure 1). By design, our approach eliminates  
 117 this requirement and enables immediate inference as soon as the model is queried.

## 119 2 BACKGROUND

120 Consider a sequence of pairs  $(c_1, y_1), (c_2, y_2), \dots, (c_n, y_n)$ , each within the measurable space  
 121  $(\mathcal{C} \times \mathcal{Y}, \mathcal{A})$ , where  $c_i$  represents the spatiotemporal coordinate,  $y_i$  denotes the corresponding solution,  
 122  $n$  is data size, and  $\mathcal{A}$  denotes the Borel  $\sigma$ -algebra on the measurable space  $\mathcal{C} \times \mathcal{Y}$ . These pairs  
 123 are drawn from a family of probability density distributions  $\{p_q : q \in \mathcal{Q}\}$ , where  $\mathcal{Q}$  represents  
 124 the *parameter space*, equipped with a  $\sigma$ -algebra  $\mathcal{B}$ , ensuring that the mappings  $q \mapsto p_q(c, y)$  are  
 125 measurable. The true parameter  $\pi$  is an element of  $\mathcal{Q}$ , and the pairs  $(c_i, y_i)$  are sampled according to  
 126  $p_\pi$ . Lacking information on  $\pi$ , we adopt a Bayesian framework to establish a prior distribution  $\Pi$   
 127 which is defined as a probability measure on  $(\mathcal{Q}, \mathcal{B})$ . Then we have

$$128 \quad \Pi(A | c, y) = \frac{\int_A p_q(c, y) d\Pi(q)}{\int_{\mathcal{Q}} p_q(c, y) d\Pi(q)},$$

131 for any measurable set  $A \in \mathcal{B}$ . This prior is updated with the observed data to form the posterior  
 132 distribution, which is defined as

$$133 \quad \Pi(A | D_n) = \frac{\int_A L_n(q) d\Pi(q)}{\int_{\mathcal{Q}} L_n(q) d\Pi(q)},$$

136 where  $L_n(q) = \prod_{i=1}^n \frac{p_q(c_i, y_i)}{p_\pi(c_i, y_i)}$  for  $A \subset \mathcal{Q}$  and  $D_n = \{(c_i, y_i)\}_{i=1}^n$ . The resulting posterior is

$$138 \quad q_n(c, y | D_n) = \int_{\mathcal{Q}} p_q(c, y) d\Pi(q | D_n),$$

140 and the posterior predictive distribution (PPD) is formulated as

$$141 \quad q_n(y | c, D_n) = \int_{\mathcal{Q}} p_q(y | c) d\Pi(q | D_n).$$

144 As noted by (Nagler, 2023; Walker, 2003; 2004b;a; Blasi & Walker, 2013), for a well-behaved prior,  
 145 the PPD converges toward true  $\pi$  as the data size  $n$  increases. This aligns with findings by Blasi &  
 146 Walker (2013), demonstrating that in well-specified scenarios, strong consistency is achieved as

$$147 \quad \Pi^n(\{q : H(p_\pi, p_q) > \epsilon\}) \rightarrow 0 \quad \text{almost surely}, \quad (1)$$

148 for any  $\epsilon > 0$ , where  $\Pi^n(A) = \int_A d\Pi(q | D_n)$  is the posterior measure and  $H$  is the Hellinger  
 149 distance defined by

$$150 \quad H(p, q) = \left( \int_{\mathcal{C} \times \mathcal{Y}} (\sqrt{p} - \sqrt{q})^2 \right)^{1/2}.$$

153 **Theorem 2.1.** Let  $D_n$ , whose size is  $n$ , be a set of ground-truth prior data, and  $\tilde{D}_n = D_n + \eta_n$  be  
 154 our (unbiased) observation, where  $\eta_n$  is a zero-mean noise distribution with a finite variance. Let  
 155  $p_\pi(\cdot | c)$  denote the true posterior, and let  $\hat{p}_\theta(\cdot | c, \tilde{D}_n)$  be the corresponding learned (approximate)  
 156 posterior for some neural network parameters  $\theta$ . Suppose that the same conditions as in Lemma A.1  
 157 hold (with  $\tilde{D}$  in place of  $D_n$ ). Then, it holds that

$$158 \quad \lim_{n \rightarrow \infty} \mathbb{E}_c \left[ H \left( \hat{p}_\theta(\cdot | c, \tilde{D}_n), p_\pi(\cdot | c) \right) \right] = 0 \quad \text{almost surely},$$

160 where  $H(\cdot, \cdot)$  denotes the Hellinger distance (see Appendix A for proof).

161 <sup>1</sup>In ICON and ICON-LM, a demo refers to a set of (input, output) pairs of an operator to be inferred.

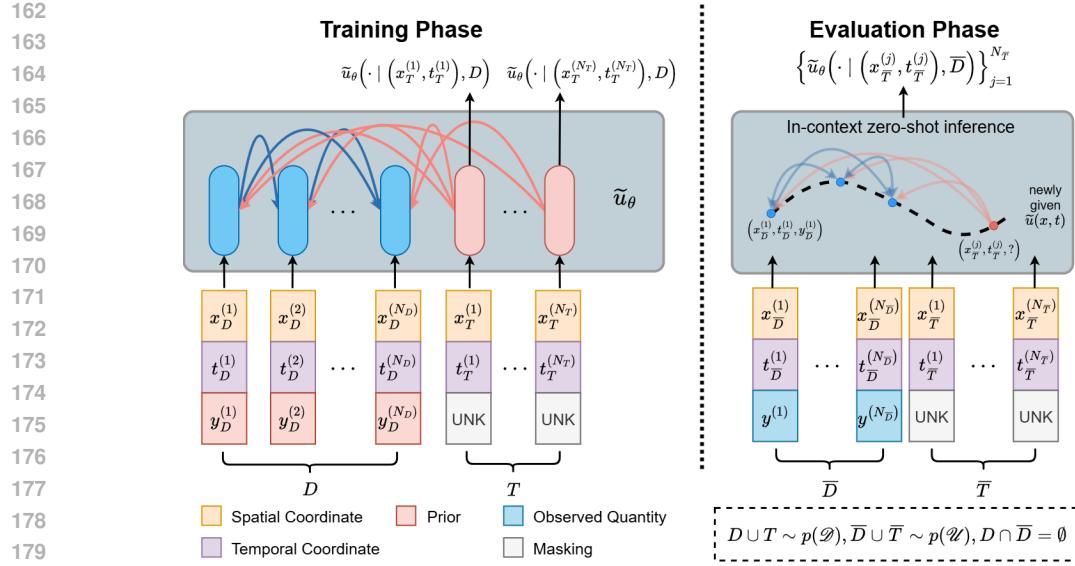


Figure 2: **A schematic diagram of our model.** (Left) Given the contexts  $D$  and queries  $T$  drawn from the prior distribution  $\mathcal{D}$ , our model  $\tilde{u}_\theta$  is trained to infer the solutions of the queried points  $T$  in the training phase. ICL is leveraged with self-attention (blue) within  $D$  and cross-attention (red) from  $T$  to  $D$ . (Right) In the testing phase, the pre-trained  $\tilde{u}_\theta$  takes contexts  $\bar{D}$  and queries  $\bar{T}$  drawn from the ground truth distribution  $\mathcal{U}$ , and predicts for the queried points  $\bar{T}$ .

This result demonstrates the sensitivity of the posterior distribution approximation, accomplished by the neural network, to perturbations of the data by zero-mean noise. As the dataset size  $n$  increases, the network becomes increasingly sensitive to the posterior distribution, converging to the expected value under the prior distribution. This sensitivity to the data reflects the consistency and robustness of the Bayesian inference process. Our model leverages this observation by performing Bayesian inference that incorporates (noisy) prior data, allowing it to infer robust solutions for given spatiotemporal conditions. The experimental results in Appendix H.4 confirm this behavior, showing how the network's solution converges to the true solution as  $n$  increases.

### 3 METHODS

In the Background section, we denoted the spatiotemporal coordinate as  $c$  for integration. However, for more clarity, as shown in Figure 2, we separate the spatiotemporal coordinate  $c$  into a spatial coordinate  $x$  and a temporal coordinate  $t$ , i.e.,  $c = (x, t)$ . Let  $\alpha$  denote a parameter vector representing the coefficients of the governing PDE dynamics. Figure 2 illustrates the schematic diagram of our model. As defined in Figure 2,  $D = \{(x_D^{(i)}, t_D^{(i)}, y_D^{(i)})\}_{i=1}^{N_D}$  and  $T = \{(x_T^{(i)}, t_T^{(i)}, y_T^{(i)})\}_{i=1}^{N_T}$  represent the context set and the query set for ICL, respectively. We assume these sets are independently and identically sampled from a family of probability density distributions  $\{p_\alpha : \alpha \in \mathcal{Q}\}$ . Here, each  $y$  corresponds to the quantity at  $(x, t)$  for a PDE solution  $u(\alpha)$  determined by the PDE parameter  $\alpha$ . Depending on the problem setting,  $u$  can be either the ground-truth solution or an approximated solution generated by a PINN, since we consider both exact and approximated prior data. The PPD of the solutions given the context set can be expressed as

$$\pi(y | x, t, D) = \int_{\mathcal{Q}} p_\alpha(y | x, t) d\Pi(\alpha | D).$$

This represents the posterior distribution of  $y$  given  $D$ , capturing the most probable solution distribution for the parameter  $\alpha$ . In this work, we aim to predict the solution from  $D$  by minimizing the mean squared error (MSE) between the PPD-derived solution and the true solution, even with noise.

**Model architecture** Since our model is essentially a prior-data fitted network (PFN), we construct its architecture by modifying certain components of the original PFN. To better handle the more complex PDE data, we first incorporate a Fourier feature embedding into the input data before it is passed to the encoder, extending the original design that was limited to tabular data. We also introduce

several modifications to the Transformer architecture. Specifically, to improve the model’s ability to learn from complex data while reducing computational complexity and enhancing representational capacity, we adopt the attentive graph filter (AGF) Wi et al. (2025) structure for Transformer layers. A more detailed description of the model architecture can be found in Appendix E.

**Benchmark PDEs** We consider convection-diffusion-reaction (CDR) equations as benchmark PDEs for diverse tasks<sup>2</sup>:

$$u_t = \mathcal{N}(\cdot), \quad \mathcal{N}(t, x, u, \beta, \nu, \rho_1, \dots, \rho_J) = -\beta \cdot \nabla u + \nabla \cdot (\nu \nabla u) + \sum_{j=1}^J \rho_j f_j(u), \quad (2)$$

where  $u$  denotes the state variable;  $\beta$  is the convection (advection) velocity;  $\nu$  denotes the anisotropic diffusivity coefficient; and  $\rho_1, \rho_2, \rho_3$  are scalar coefficients corresponding to the nonlinear reaction terms  $f_1, f_2, f_3 : \mathbb{R} \rightarrow \mathbb{R}$ , respectively. For clarity, we will refer to the two-dimensional CDR equations, including cases where certain coefficients are zero, as the family of CDR equations. This family encompasses not only generic parabolic equations but also extends to hyperbolic equations such as convection equations. Hence, the family of CDR equations contains three key terms—convection, diffusion, and reaction—each with distinct properties, making it an ideal benchmark problem. To the best of our knowledge, this work is the first to address a wide range of reaction terms within the family of CDR equations using a single unified model. We use a set of CDR-related terms and a linear combination of three nonlinear reaction terms to generate prior data. This formulation enables the incorporation of diverse reaction dynamics, making the benchmark problems more comprehensive.

**PINN-based approximations of PDE solution space** Although numerical solvers can solve the family of CDR equations, in this work, we use the predictions from PINNs as a representative example of low-cost approximated data. This allows us to evaluate the robustness of our model to noise, which is commonly encountered in real-world scenarios. From this point, we refer to these PINN-based approximations as PINN priors. To approximate the solution space for the PDEs, we construct a practical parameter subspace,  $\Omega \subseteq \mathcal{Q}$ , which is the collection of coefficients in Eq. (2) and has a dictionary form  $\Omega = \{\alpha := (\beta, \nu, \rho_1, \dots, \rho_J)\}$ . Consequently, the target exact prior  $\mathcal{U}$  defined below represents the collection of solutions  $u(\alpha)$  at Eq. (2) for each parameter  $\alpha \in \Omega$ , where  $\mathcal{X}$  and  $\mathcal{T}$  correspond to the spatial and temporal domains of interest, respectively.

$$\mathcal{U} = \bigcup_{\alpha \in \Omega} \{u(\alpha) \mid u_t = \mathcal{N}(t, x, u, \alpha)\}, \quad u : \mathcal{X} \times \mathcal{T} \rightarrow \mathbb{R}.$$

Since the target exact prior data  $\mathcal{U}$  is hard to obtain, we instead use a PINN prior  $\mathcal{D}$  that closely approximates  $\mathcal{U}$  as follows. Suppose  $\tilde{u}(\alpha)$  is the prediction by PINN, which is trained to predict the PDE  $u_t = \mathcal{N}(\cdot)$  (see Appendix C for details). The PINN prior  $\mathcal{D}$  is a collection of approximate solutions  $\tilde{u}(\alpha)$  for each  $\alpha \in \Omega$ ,

$$\mathcal{D} = \bigcup_{\alpha \in \Omega} \{\tilde{u}(\alpha)\}, \quad p(\mathcal{D}) \sim p(\mathcal{U}).$$

Subsequently, the model performs the PPD of the generated prior  $p(\mathcal{D})$  through ICL.

**Objective function** From a given parameter space  $\Omega$ , the parameter  $\alpha$  is randomly drawn i.i.d. from  $\Omega$ . After this, the previous  $\tilde{u}(\alpha)$  is given as an input to our model  $\tilde{u}_\theta$  to minimize the MSE at the predicted points, as expressed in Eq. (3). The specific MSE loss chosen to regress the solution over the spatial and temporal domain for a given  $\tilde{u}(\alpha)$  is defined as

$$\mathcal{L}_\alpha = \frac{1}{N_T} \sum_{j=1}^{N_T} \left[ \tilde{u}_\theta(x_T^{(j)}, t_T^{(j)} \mid D) - \tilde{u}(x_T^{(j)}, t_T^{(j)}) \right]^2. \quad (3)$$

In addition, we introduce a regularization term into the objective function to ensure orthogonality during the SVD process within the AGF layer (see Appendix E for detail). We denote this term as

<sup>2</sup>We note that in Eq. (2)  $J$  reaction terms are considered. In previous works (Cho et al., 2023; 2024; Kang et al., 2024), only one type of reaction term is considered. Therefore, one can consider that Eq. (2) is a comprehensive dictionary (or a prior set for PFN training) of all popular CDR equations. Our experiments differ from previous works in this regard.

270  $\mathcal{L}_{AGF}$ . To balance the trade-off between the loss and the regularization, we scale the latter by the  
 271 **number of context points  $N_D$** . Consequently, the final objective function  $\mathcal{L}$  is expressed as  
 272

$$273 \quad \mathcal{L} = \mathcal{L}_\alpha + \frac{1}{N_D} \mathcal{L}_{AGF}. \\ 274$$

275 The concrete flow of the training and the evaluation process is described in Appendix F and G.  
 276

## 277 4 EXPERIMENTS

278 Our experimental section is divided into three phases. In the first phase, we introduce the dataset  
 279 and establish the baselines for our experiments. In the second phase, we pre-train our model on  
 280 the family of 2D CDR equations and evaluate it on a variety of tasks without any additional task-  
 281 specific fine-tuning. This allows us to examine the model’s ability to handle diverse tasks. In the  
 282 third phase, we fine-tune our model on different PDEs, including the shallow water equations (SWE)  
 283 and compressible Navier-Stokes equations (CNSE), in order to assess its generalization performance  
 284 across different PDEs. Table 1 summarizes the experiments in this section.

285 Table 1: Roadmap of the experiment section. ‘Inter’ and ‘Extra’ denote spatiotemporal interpolation  
 286 and temporal extrapolation tasks, respectively. The ‘Coeff/dataset’ column specifies either the  
 287 coefficients of the equation or the dataset used in evaluation.

289 Section	290 Task	291 Coeff/Dataset	292 Equation	293 Objective
4.2.1	Inter & Extra	Seen	Family of 2D CDR equations	Task generalization
4.2.2	Inter & Extra	Unseen		Task generalization in parameter space
4.3.1	Inter & Extra	Seen & Unseen	SWE	PDE generalization
4.3.2	Operator Learning	Unseen	CNSE	PDE generalization for new task

### 4.1 EXPERIMENTAL SETUP

294 **Dataset** We consider a family of 2D CDR equations with three reaction terms as the  
 295 pre-training dataset. This family is constructed from equations with coefficient dictionary  
 296 ( $\beta_x, \beta_y, \nu_x, \nu_y, \rho_1, \rho_2, \rho_3 \in \{0.0, 0.5, 1.0\}$ ), resulting in a total of 2,187 unique PDEs. **All PDE**  
 297 **instances are simulated from the same initial condition, which allows differences in the resulting**  
 298 **dynamics to stem solely from the variations in coefficients.** For PINN prior, each PINN is trained  
 299 independently, and its predicted solutions are used as noisy priors. On average, the generated PINN  
 300 priors exhibit an  $L_2$  relative error of 0.02280. For PDE generalization experiments, we employ the  
 301 SWE and CNSE datasets from PDEBench (Takamoto et al., 2022). These two equations are chosen  
 302 because they are commonly used to validate PDE solvers in SciML. Both datasets consist of 1,000  
 303 samples generated from distinct initial conditions. The SWE dataset contains a single feature, the  
 304 water height, whereas the CNSE dataset includes  $x$ - and  $y$ -velocities, pressure, and density. Detailed  
 305 information on the datasets can be found in Appendix B.

306 **Baseline methods** We compare our model with six baseline models: deep operator network (Deep-  
 307 ONet) (Lu et al., 2021), Fourier neural operator (FNO) (Li et al., 2020), factorized Fourier neural  
 308 operator (F-FNO) (Tran et al., 2021), adaptive Fourier neural operator (A-FNO) (Guibas et al.,  
 309 2021), **latent neural operator (LNO)** (Wang & Wang, 2024), Poseidon (Herde et al., 2024), and auto-  
 310 regressive denoising pre-training operator Transformer (DPOT) (Hao et al., 2024). These baselines  
 311 are chosen for their scalability and ability to effectively handle complex 2D PDEs. For Poseidon and  
 312 DPOT, scientific foundation models with publicly available pre-trained weights, we use the released  
 313 weights and perform fine-tuning. The other baselines are trained and evaluated separately on each  
 314 dataset and task. For the spatiotemporal interpolation task, DeepONet is the only baseline used,  
 315 excluding those that require grid inputs. The best hyperparameters for all models are selected based  
 316 on their best validation performance in each task. Detailed information on the baselines used in each  
 317 experiment can be found in Appendix D.

### 4.2 PRE-TRAINING ON THE FAMILY OF 2D CDR EQUATIONS AND EVALUATION

318 The training dataset is constructed from 6 time steps within the range  $t \in [0.0, 0.5]$  at intervals of  
 319 0.1. Our model takes as an input set of contexts, which include domain coordinates paired with their  
 320 solution values, along with queries, which consist of domain coordinates where the solution values  
 321 are masked. We denote the version of our model pre-trained on numeric priors as **Ours** and the  
 322 version pre-trained on PINN priors as **Ours (PINN)**. For the baselines, the training procedure differs  
 323

324 by task. In the spatiotemporal interpolation task, since DeepONet requires fixed points as branch  
 325 input, we predefined the fixed points in training data by sampling and the remaining points are used  
 326 as queries to be predicted. In the temporal extrapolation task, the baselines are trained to take the  
 327 grid points at a time  $t_n$  as input and predict the grid points at the subsequent time step  $t_{n+1}$ . Further  
 328 information on the training procedure, such as the number of epochs, can be found in Appendix I.  
 329

#### 330 4.2.1 TASK GENERALIZATION FOR SEEN COEFFICIENTS

331 We first evaluate the task generalization ability in spatiotemporal interpolation and temporal extrapolation  
 332 without additional task-specific fine-tuning. Both tasks are conducted within the coefficient  
 333 dictionary used during pre-training, which we refer to as seen coefficients. For the spatiotemporal  
 334 interpolation task, the baseline DeepONet requires fixed points as branch input. Therefore, our model  
 335 is also provided with the same set of context points during evaluation to ensure a fair comparison.  
 336 Specifically, we use the points at  $t = 0.25$  as validation queries and the points at  $t = 0.05, 0.15, 0.35$ ,  
 337 and  $0.45$  as test queries. For the temporal extrapolation task, the points at  $t = 0.5$  are used as contexts,  
 338 and those at  $t = 0.6$  as queries for validation. At test time, the points at  $t = 0.6$  serve as contexts,  
 339 and those at  $t = 0.7, 0.8, 0.9$ , and  $1.0$  are used as queries. The baselines predict a sequence of query  
 340 points through rollout, whereas our model can directly predict the query points from their coordinates  
 341 so that it is evaluated in this manner. The evaluation results are shown in Table 2.  
 342

343 Table 2: The evaluation results for the spatiotemporal interpolation and temporal extrapolation  
 344 tasks applied to the family of 2D CDR equations. They are measured at the seen coefficients  
 345 ( $\beta_x, \beta_y, \nu_x, \nu_y, \rho_1, \rho_2, \rho_3 \in \{0.0, 0.5, 1.0\}$ ). The best performance is marked in **bold** and the second-  
 346 best performance is marked with underline.

Task	Metric	Ours	Ours (PINN)	DeepONet	FNO	F-FNO	A-FNO	LNO	Poseidon	DPOT
Spatiotemporal interpolation	$L_1$ Abs	<b>0.01218</b>	<u>0.01507</u>	0.01922	-	-	-	-	-	-
	$L_2$ Rel	<b>0.01643</b>	<u>0.02057</u>	0.02479	-	-	-	-	-	-
	$L_\infty$ Rel	<b>0.04600</b>	0.06584	<u>0.05980</u>	-	-	-	-	-	-
Temporal extrapolation	$L_1$ Abs	<b>0.01474</b>	<u>0.01940</u>	0.06529	0.03302	0.02731	0.05800	<b>0.02876</b>	0.14252	0.04924
	$L_2$ Rel	<b>0.02261</b>	<u>0.02705</u>	0.08212	0.04058	0.03430	0.07915	<b>0.03665</b>	0.15580	0.08587
	$L_\infty$ Rel	<b>0.08444</b>	0.08802	0.22802	0.09553	0.09877	0.28905	<b>0.10948</b>	0.28285	<u>0.08587</u>

352 The evaluation results demonstrate that both versions of our model outperform the baselines on both  
 353 tasks. This indicates that once trained on a single dataset, our model can achieve strong performance  
 354 without requiring task-specific fine-tuning. Even when trained with noisy PINN priors, our model  
 355 surpasses the baselines in two metrics, confirming its robustness to noise in the priors. Additionally,  
 356 we conduct experiments where the PINN prior is provided as a context set during evaluation in  
 357 Appendix H.1.  
 358

#### 4.2.2 TASK GENERALIZATION FOR UNSEEN COEFFICIENTS

359 We further evaluate the task generalization ability on both tasks in parameter space using the family  
 360 of 2D CDR equations with unseen coefficients, without any additional fine-tuning. The unseen  
 361 coefficients are categorized into two cases: coefficient interpolation and coefficient extrapolation.  
 362 Coefficient interpolation experiments are performed on the equations with intermediate coefficient  
 363 values not seen during pre-training, ( $\beta_x, \beta_y, \nu_x, \nu_y, \rho_1, \rho_2, \rho_3 \in \{0.25, 0.75\}$ ). Coefficient extrapolation  
 364 experiments, on the other hand, are conducted on the equations with larger coefficient values  
 365 than those used in pre-training, ( $\beta_x, \beta_y, \nu_x, \nu_y, \rho_1, \rho_2, \rho_3 \in \{1.25, 1.5\}$ ). For evaluation with unseen  
 366 coefficients in the temporal extrapolation task, we modify the temporal extrapolation setup to remain  
 367 within the temporal domain used during training: the points at  $t = 0.1$  are provided as contexts,  
 368 and models are asked to predict the values at  $t = 0.2, 0.3, 0.4$ , and  $0.5$ . The evaluation results are  
 369 summarized in Table 3.

370 The evaluation results show that both versions of our model outperform the baselines on both tasks  
 371 under coefficient interpolation and extrapolation. This shows that our models can make accurate  
 372 predictions not only on PDEs with coefficients seen during training but also on PDEs with unseen  
 373 coefficients, without requiring additional fine-tuning. In other words, the models are able to generalize  
 374 across the parameter space, achieving both interpolation and extrapolation.

#### 375 4.3 FINE-TUNING ON DIFFERENT PDES AND EVALUATION

376 To evaluate the PDE generalization ability of our model, we conduct experiments on heterogeneous  
 377 PDEs beyond the family of 2D CDR equations. On the shallow water equations (SWE) dataset, we

378  
379  
380  
381  
382  
383  
384  
385  
386  
387  
388  
389  
390  
391  
392  
393  
394  
395  
396  
397  
398  
399  
400  
401  
402  
403  
404  
405  
406  
407  
408  
409  
410  
411  
412  
413  
414  
415  
416  
417  
418  
419  
420  
421  
422  
423  
424  
425  
426  
427  
428  
429  
430  
431

Table 3: The evaluation results for the spatiotemporal interpolation and temporal extrapolation tasks applied to the family of 2D CDR equations. They are measured at the unseen coefficients. The best performance is marked in **bold** and the second-best performance is marked with underline.

Coefficients	Task	Metric	Ours	Ours (PINN)	DeepONet	FNO	F-FNO	A-FNO	LNO	Poseidon	DPOT
Interpolation	Spatiotemporal interpolation	$L_1$ Abs	<b>0.01241</b>	<u>0.01277</u>	0.01347	-	-	-	-	-	-
		$L_2$ Rel	<b>0.01658</b>	0.01761	0.01731	-	-	-	-	-	-
		$L_\infty$ Rel	0.04274	0.05503	<b>0.04114</b>	-	-	-	-	-	-
	Temporal extrapolation	$L_1$ Abs	<u>0.08957</u>	<b>0.08709</b>	0.11918	0.11625	0.11194	0.12220	<b>0.11085</b>	0.17740	0.13879
		$L_2$ Rel	<b>0.11990</b>	<u>0.12465</u>	0.14903	0.14658	0.14348	0.18191	<b>0.14913</b>	0.21340	0.15085
		$L_\infty$ Rel	0.29526	0.36472	0.35196	<b>0.28689</b>	0.30813	0.64429	<b>0.33280</b>	0.39890	0.15085
Extrapolation	Spatiotemporal interpolation	$L_1$ Abs	0.01661	<b>0.01582</b>	0.03576	-	-	-	-	-	-
		$L_2$ Rel	0.02398	<b>0.02038</b>	0.04669	-	-	-	-	-	-
		$L_\infty$ Rel	0.08324	<b>0.06890</b>	0.10392	-	-	-	-	-	-
	Temporal extrapolation	$L_1$ Abs	<b>0.01027</b>	<u>0.02063</u>	0.05514	0.02652	0.02854	0.05013	<b>0.05076</b>	0.14661	0.05186
		$L_2$ Rel	<b>0.01068</b>	0.02169	0.06450	0.03043	0.03300	0.06351	<b>0.06347</b>	0.15650	0.05189
		$L_\infty$ Rel	<b>0.01661</b>	<u>0.03474</u>	0.16555	0.06124	0.08343	0.20955	<b>0.15091</b>	0.27502	0.05189

assessed whether our model could generalize, via fine-tuning, to heterogeneous PDEs on the same tasks as in the family, namely spatiotemporal interpolation and temporal extrapolation. The baselines for each task are trained in the same manner as in the family. On the CNSE dataset, we further examined whether the model could generalize to a new task, operator learning, under a different PDE.

#### 4.3.1 PDE GENERALIZATION FOR SHALLOW WATER EQUATIONS

To evaluate the PDE generalization ability of our model, we conduct experiments on the SWE, assessing whether the tasks defined in the family of 2D CDR equations can be successfully transferred to this heterogeneous PDE by minimal fine-tuning. Out of the total 1,000 samples, 700 are used for training, and the training, validation, and test datasets are constructed over the time interval in the same manner as for the family of 2D CDR equations. The test dataset generated from the initial condition samples used during training is referred to as the seen dataset, while the test dataset generated from the remaining 300 samples is referred to as the unseen dataset. For the baselines, training for each task is carried out in the same manner as in pre-training.

Table 4: The evaluation results for the spatiotemporal interpolation and temporal extrapolation tasks applied to SWE. They are measured at the seen/unseen datasets. The best performance is marked in **bold** and the second-best performance is marked with underline.

Evaluation dataset	Task	Metric	Ours	Ours (PINN)	DeepONet	FNO	F-FNO	A-FNO	LNO	Poseidon	DPOT
Seen	Spatiotemporal interpolation	$L_1$ Abs	<b>0.00208</b>	0.00269	0.00271	-	-	-	-	-	-
		$L_2$ Rel	<b>0.00739</b>	<u>0.00822</u>	0.00997	-	-	-	-	-	-
		$L_\infty$ Rel	<b>0.05897</b>	0.06950	0.06951	-	-	-	-	-	-
	Temporal extrapolation	$L_1$ Abs	<u>0.01548</u>	<b>0.01418</b>	0.06994	0.03271	0.02292	0.05036	<b>0.02902</b>	0.02964	0.03327
		$L_2$ Rel	<b>0.03273</b>	<u>0.03221</u>	0.08319	0.05429	0.04110	0.07436	<b>0.05269</b>	0.05208	0.05860
		$L_\infty$ Rel	<b>0.20590</b>	0.24596	0.28736	0.24954	<u>0.24478</u>	0.38030	<b>0.24832</b>	0.25668	0.29727
Unseen	Spatiotemporal interpolation	$L_1$ Abs	<b>0.00213</b>	<u>0.00272</u>	0.00278	-	-	-	-	-	-
		$L_2$ Rel	<b>0.00754</b>	<u>0.00828</u>	0.01008	-	-	-	-	-	-
		$L_\infty$ Rel	<b>0.05986</b>	<u>0.06981</u>	0.07047	-	-	-	-	-	-
	Temporal extrapolation	$L_1$ Abs	<u>0.01571</u>	<b>0.01448</b>	0.07047	0.03342	0.02371	0.05099	<b>0.02877</b>	0.03021	0.03372
		$L_2$ Rel	<b>0.03323</b>	<u>0.03288</u>	0.08347	0.05530	0.04218	0.07513	<b>0.05225</b>	0.05306	0.05914
		$L_\infty$ Rel	<b>0.20686</b>	0.24686	0.28441	0.24980	<u>0.24378</u>	0.38196	<b>0.24788</b>	0.25710	0.29757

The experimental results in Table 4 show that both versions of our model outperform the baselines on both spatiotemporal interpolation and temporal extrapolation tasks, regardless of whether the dataset is seen or unseen. Specifically, Ours achieves the best performance on spatiotemporal interpolation, while Ours (PINN) performs best on temporal extrapolation. These results demonstrate that our approach can generalize the two tasks verified in pre-training to different PDEs through fine-tuning, while also retaining generalization ability to unseen datasets.

#### 4.3.2 PDE GENERALIZATION FOR COMPRESSIBLE NAVIER-STOKES EQUATIONS

To further examine the PDE generalization ability of our model, we conduct experiments on the compressible Navier-Stokes equations (CNSE), focusing in particular on learning the time trajectory predicting operator. This experiment is designed not only to verify generalization across heterogeneous PDEs but also to assess the potential of our approach in learning an operator that was not included in previous experiments. The dataset consists of 1,000 samples, of which 700 are used for training, 100 for validation, and 200 for testing. The target operator is defined as follows: given the initial condition of four features, the models are required to predict the solution of them at later times

*t* = 0.25, 0.5, 0.75, and 1.0 (see Appendix B.3 for a formal description). Baseline models except DeepONet and DPOT are trained and evaluated using rollout predictions. In contrast, DeepONet and our model directly predict the solutions given the initial condition. Additionally, since DPOT allows the output timestep length to be adjusted, it also predicts the solutions in a single step.

Table 5: The evaluation results for the operator learning task applied to CNSE. The best performance is marked in **bold** and the second-best performance is marked with underline.

Task	Metric	Ours	Ours (PINN)	DeepONet	FNO	F-FNO	A-FNO	LNO	Poseidon	DPOT
Operator learning	$L_1$ Abs	0.09497	0.09429	0.17797	0.09581	<b>0.08734</b>	0.13637	<u>0.10366</u>	0.13313	<u>0.09164</u>
	$L_2$ Rel	<u>0.54833</u>	<b>0.54369</b>	1.09567	0.57905	0.55420	0.65798	<u>0.62536</u>	0.66904	0.67869
	$L_\infty$ Rel	<u>0.63533</u>	<b>0.63374</b>	1.09165	0.70708	0.72876	0.85729	<u>0.66936</u>	0.92909	1.15028

The experimental results in Table 5 show that both versions of our model outperform the baselines on two relative error metrics. Notably, although DPOT’s pre-training corpus already includes the CNSE dataset from PDEBench, our models still achieve superior performance over DPOT on both metrics. While DPOT was originally evaluated on predicting one time step ahead, in our setting the task was extended to predicting up to four time steps ahead. As a result, DPOT exhibited lower performance than reported in the original paper. Experimental results confirm the generalization capability of our approach for operator learning on heterogeneous PDEs involving multiple variables.

To further verify PDE generalization and input flexibility, we conduct supplementary experiments. First, we evaluate operator learning in the same form as in the CNSE experiments on the Airfoil dataset, which is constructed with an irregular mesh structure. In addition, we evaluate operator learning across different features, rather than within the same features, using the Darcy Flow dataset. Detailed information on these experiments is provided in Appendix H.

## 5 RELATED WORKS

**In-context learning** Transformers have shown remarkable ICL abilities in various studies. They generalize to unseen tasks by emulating Bayesian predictors (Panwar et al., 2024) and linear models (Zhang et al., 2024), while also efficiently performing Bayesian inference through Prior-Data Fitted Networks (PFNs) (Müller et al., 2021). Their robustness extends to learning different function classes, such as linear and sparse linear functions, decision trees, and two-layer neural networks, even under distribution shifts (Garg et al., 2022). Furthermore, Transformers can adaptively select algorithms based on input sequences, achieving near-optimal performance on tasks like noisy linear models (Bai et al., 2023). They are also fast and effective for tabular data classification (Hollmann et al., 2022).

**Scientific foundation models** Recent studies have advanced in-context operator learning and PDE solving through Transformer-based models. The work Ye et al. (2024) introduces PDEformer, a versatile model for solving 1D PDEs with high accuracy in inverse problems. In-context operator learning has also been extended to multi-modal frameworks, as seen in Yang et al. (2025), where ICON-LM integrates natural language with equations to outperform traditional models. Additionally, Yang & Osher (2024) and Yang et al. (2023) demonstrate the generalization capabilities of ICON in solving various PDE-related tasks, highlighting ICON’s few-shot learning performance across various problems in differential equations. Several other studies have addressed the problem of solving diverse PDEs using a single trained model (Hang et al., 2024; McCabe et al., 2023; Herde et al., 2024; Hao et al., 2024). However, many of these approaches rely on symbolic PDE information, true or near-true solutions, and/or do not support zero-shot, in-context learning, making their objectives different from ours.

## 6 CONCLUSION

We presented PDE-PFN, a new method for scientific machine learning that integrates in-context learning and Bayesian inference to directly approximate the posterior predictive distribution of PDE solutions. Our experiments demonstrated that PDE-PFN achieves both task and PDE generalization, handles flexible input structures without relying on governing equations, and remains robust even under noisy priors, while also enabling zero-shot inference. Together, these properties establish PDE-PFN as a flexible and robust foundation for advancing scientific machine learning. Since our current study focuses on demonstrating the feasibility of a PFN in SciML, its empirical verification is limited to two-dimensional PDEs and has not yet been validated on higher-dimensional or more complex systems. Future work will extend PDE-PFN to such challenging settings, further enhancing its applicability to real-world scientific and engineering problems.

486  
487  
**ETHICS STATEMENT**488  
489  
490  
491  
492  
493  
This research adheres to the ethical standards required for scientific inquiry. We have considered  
the potential societal impacts of our work and have found no clear negative implications. We also  
see no direct path from our model to malicious uses such as surveillance, disinformation, or privacy  
breaches. On the other hand, we recognize that improvements in PDE-solving capabilities could  
eventually be applied in real-world applications, such as climate modeling, biomedical simulations,  
or engineering systems.494  
495  
496  
All experiments were conducted in compliance with relevant laws and ethical guidelines, ensuring  
the integrity of our findings. We are committed to transparency and reproducibility in our research  
processes.497  
498  
**REPRODUCIBILITY**  
499500  
501  
502  
503  
We are committed to ensuring the reproducibility of our research. All experimental procedures, data  
sources, and algorithms used in this study are clearly documented in the paper. The code will be  
provided as the supplementary material and be made publicly available upon publication, allowing  
others to validate our findings and build upon our work.504  
505  
**REFERENCES**  
506507  
508  
509  
510  
Yu Bai, Fan Chen, Huan Wang, Caiming Xiong, and Song Mei. Transformers as statisticians: Provable  
in-context learning with in-context algorithm selection. In A. Oh, T. Naumann, A. Globerson,  
K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*,  
volume 36, pp. 57125–57211. Curran Associates, Inc., 2023.511  
512  
Andrea Beck and Marius Kurz. A perspective on machine learning methods in turbulence modeling.  
*GAMM-Mitteilungen*, 44(1):e202100002, 2021.513  
514  
Pierpaolo De Blasi and Stephen G. Walker. Bayesian asymptotics with misspecified models. *Statistica  
Sinica*, pp. 169–187, 2013.516  
517  
Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal,  
Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel  
Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler,  
Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray,  
Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever,  
and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato,  
R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*,  
volume 33, pp. 1877–1901. Curran Associates, Inc., 2020.524  
525  
Chen-Fu Chien, Chia-Yu Hsu, and Chih-Wei Hsiao. Manufacturing intelligence to forecast and  
reduce semiconductor cycle time. *Journal of Intelligent Manufacturing*, 23:2281–2294, 2012.526  
527  
Woojin Cho, Kookjin Lee, Donsub Rim, and Noseong Park. Hypernetwork-based meta-learning for  
low-rank physics-informed neural networks. In A. Oh, T. Naumann, A. Globerson, K. Saenko,  
M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36,  
pp. 11219–11231. Curran Associates, Inc., 2023.531  
532  
Woojin Cho, Minju Jo, Haksoo Lim, Kookjin Lee, Dongeun Lee, Sanghyun Hong, and Noseong Park.  
Parameterized physics-informed neural networks for parameterized pdes. In *Proceedings of the  
41st International Conference on Machine Learning*, pp. 8510–8533, 2024.534  
535  
Junho Choi, Taehyun Yun, Namjung Kim, and Youngjoon Hong. Spectral operator learning for  
parametric pdes without data reliance. *Computer Methods in Applied Mechanics and Engineering*,  
420:116678, 2024. ISSN 0045-7825. doi: <https://doi.org/10.1016/j.cma.2023.116678>.538  
539  
Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam  
Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh,  
Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam

540 Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James  
 541 Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Lev-  
 542 skaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin  
 543 Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph,  
 544 Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M.  
 545 Dai, Thanumalayan Sankaranarayana Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon  
 546 Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark  
 547 Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean,  
 548 Slav Petrov, and Noah Fiedel. Palm: Scaling language modeling with pathways. *Journal of  
 549 Machine Learning Research*, 24(240):1–113, 2023.

550 Simon Frieder, Luca Pinchetti, , Ryan-Rhys Griffiths, Tommaso Salvatori, Thomas Lukasiewicz,  
 551 Philipp Petersen, and Julius Berner. Mathematical capabilities of chatgpt. In A. Oh, T. Nau-  
 552 mann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information  
 553 Processing Systems*, volume 36, pp. 27699–27744. Curran Associates, Inc., 2023.

554 Shivam Garg, Dimitris Tsipras, Percy S Liang, and Gregory Valiant. What can transformers learn  
 555 in-context? a case study of simple function classes. In S. Koyejo, S. Mohamed, A. Agarwal,  
 556 D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*,  
 557 volume 35, pp. 30583–30598. Curran Associates, Inc., 2022.

558 John Guibas, Morteza Mardani, Zongyi Li, Andrew Tao, Anima Anandkumar, and Bryan Catan-  
 559 zaro. Adaptive fourier neural operators: Efficient token mixers for transformers. *arXiv preprint  
 560 arXiv:2111.13587*, 2021.

561 Zhou Hang, Yuezhou Ma, Haixu Wu, Haowen Wang, and Mingsheng Long. Unisolver: Pde-  
 562 conditional transformers are universal pde solvers. *arXiv preprint arXiv:2405.17527*, 2024.

563 Zhongkai Hao, Chang Su, Songming Liu, Julius Berner, Chengyang Ying, Hang Su, Anima Anandku-  
 564 mar, Jian Song, and Jun Zhu. Dpot: Auto-regressive denoising operator transformer for large-scale  
 565 pde pre-training. *arXiv preprint arXiv:2403.03542*, 2024.

566 Maximilian Herde, Bogdan Raonić, Tobias Rohner, Roger Käppeli, Roberto Molinaro, Emmanuel  
 567 de Bézenac, and Siddhartha Mishra. Poseidon: Efficient foundation models for pdes. *arXiv preprint  
 568 arXiv:2405.19101*, 2024.

569 Noah Hollmann, Samuel Müller, Katharina Eggensperger, and Frank Hutter. TabPFN: A transformer  
 570 that solves small tabular classification problems in a second. In *NeurIPS 2022 First Table  
 571 Representation Workshop*, 2022.

572 Mingu Kang, Dongseok Lee, Woojin Cho, Kookjin Lee, Anthony Gruber, Nathaniel Trask, Youngjoon  
 573 Hong, and Noseong Park. Can we pre-train icl-based sfms for the zero-shot inference of the 1d cdr  
 574 problem with noisy data? In *Neurips 2024 Workshop Foundation Models for Science: Progress,  
 575 Opportunities, and Challenges*, 2024.

576 Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott  
 577 Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models.  
 578 *arXiv preprint arXiv:2001.08361*, 2020.

579 David I Ketcheson, Kyle Mandli, Aron J Ahmadia, Amal Alghamdi, Manuel Quezada De Luna,  
 580 Matteo Parsani, Matthew G Knepley, and Matthew Emmett. Pyclaw: Accessible, extensible,  
 581 scalable tools for wave propagation problems. *SIAM Journal on Scientific Computing*, 34(4):  
 582 C210–C231, 2012.

583 Namjung Kim, Dongseok Lee, and Youngjoon Hong. Data-efficient deep generative model with  
 584 discrete latent representation for high-fidelity digital materials. *ACS Materials Letters*, 5(3):  
 585 730–737, 2023.

586 Namjung Kim, Dongseok Lee, Chanyoung Kim, Dosung Lee, and Youngjoon Hong. Simple arith-  
 587 metic operation in latent space can generate a novel three dimensional graph metamaterials. *arXiv  
 588 preprint arXiv:2404.06671*, 2024.

594 Chin Yik Lee and Stewart Cant. A grid-induced and physics-informed machine learning cfd frame-  
595 work for turbulent flows. *Flow, Turbulence and Combustion*, 112(2):407–442, 2024.  
596

597 Zijie Li, Kazem Meidani, and Amir Barati Farimani. Transformer for partial differential equations’  
598 operator learning. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL  
599 <https://openreview.net/forum?id=EPPqt3uERT>.

600 Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew  
601 Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations.  
602 *arXiv preprint arXiv:2010.08895*, 2020.

603 Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning  
604 nonlinear operators via deeponet based on the universal approximation theorem of operators.  
605 *Nature machine intelligence*, 3(3):218–229, 2021.

606 Mingze Ma, Qingtian Zhu, Yifan Zhan, Zhengwei Yin, Hongjun Wang, and Yinqiang Zheng.  
607 Robustifying fourier features embeddings for implicit neural representations. *arXiv preprint*  
608 *arXiv:2502.05482*, 2025.

609 Michael McCabe, Bruno Régaldo-Saint Blancard, Liam Holden Parker, Ruben Ohana, Miles Cranmer,  
610 Alberto Bietti, Michael Eickenberg, Siavash Golkar, Geraud Krawezik, Francois Lanusse, et al.  
611 Multiple physics pretraining for physical surrogate models. *arXiv preprint arXiv:2310.02994*,  
612 2023.

613 Xuhui Meng, Zhen Li, Dongkun Zhang, and George Em Karniadakis. Ppinn: Parareal physics-  
614 informed neural network for time-dependent pdes. *Computer Methods in Applied Mechanics and*  
615 *Engineering*, 370:113250, 2020.

616 Alejandro Molina, Patrick Schramowski, and Kristian Kersting. Padé activation units: End-to-end  
617 learning of flexible activation functions in deep networks. In *International Conference on Learning*  
618 *Representations*, 2019.

619 Samuel Müller, Noah Hollmann, Sebastian Pineda Arango, Josif Grabocka, and Frank Hutter. Trans-  
620 formers can do bayesian inference. In *International Conference on Learning Representations*,  
621 2022.

622 Stephen C Myers and Craig A Schultz. Improving sparse network seismic location with bayesian  
623 kriging and teleseismically constrained calibration events. *Bulletin of the Seismological Society of*  
624 *America*, 90(1):199–211, 2000.

625 Samuel Müller, Noah Hollmann, Sebastian Pineda-Arango, Josif Grabocka, and Frank Hutter. Trans-  
626 formers can do bayesian inference. *CoRR*, abs/2112.10510, 2021.

627 Thomas Nagler. Statistical foundations of prior-data fitted networks. In *International Conference on*  
628 *Machine Learning*, pp. 25660–25676. PMLR, 2023.

629 Zachary G Nicolaou, Guanyu Huo, Yihui Chen, Steven L Brunton, and J Nathan Kutz. Data-driven  
630 discovery and extrapolation of parameterized pattern-forming dynamics. *Physical Review Research*,  
631 5(4):L042017, 2023.

632 Francisco Palacios, Juan Alonso, Karthikeyan Duraisamy, Michael Colonno, Jason Hicken, Aniket  
633 Aranake, Alejandro Campos, Sean Copeland, Thomas Economou, Amrita Lonkar, et al. Stanford  
634 university unstructured (su 2): an open-source integrated computational environment for multi-  
635 physics simulation and design. In *51st AIAA aerospace sciences meeting including the new*  
636 *horizons forum and aerospace exposition*, pp. 287, 2013.

637 Madhur Panwar, Kabir Ahuja, and Navin Goyal. In-context learning through the bayesian prism. In  
638 *The Twelfth International Conference on Learning Representations*, 2024.

639 Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter Battaglia. Learning mesh-based  
640 simulation with graph networks. In *International conference on learning representations*, 2020.

641 Alfio Quarteroni and Alberto Valli. *Numerical approximation of partial differential equations*,  
642 volume 23. Springer Science & Business Media, Berlin, 2008.

648 Michael Quirk and Julian Serda. *Semiconductor manufacturing technology*, volume 1. Prentice Hall  
 649 Upper Saddle River, NJ, 2001.  
 650

651 M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning  
 652 framework for solving forward and inverse problems involving nonlinear partial differential  
 653 equations. *Journal of Computational Physics*, 378:686–707, 2019a. ISSN 0021-9991. doi:  
 654 <https://doi.org/10.1016/j.jcp.2018.10.045>.

655 Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A  
 656 deep learning framework for solving forward and inverse problems involving nonlinear partial  
 657 differential equations. *Journal of Computational physics*, 378:686–707, 2019b.

658 Nusrat Rouf, Majid Bashir Malik, Tasleem Arif, Sparsh Sharma, Saurabh Singh, Satyabrata Aich,  
 659 and Hee-Cheol Kim. Stock market prediction using machine learning techniques: a decade survey  
 660 on methodologies, recent developments, and future directions. *Electronics*, 10(21):2717, 2021.

661

662 Khemraj Shukla, Patricio Clark Di Leoni, James Blackshire, Daniel Sparkman, and George Em  
 663 Karniadakis. Physics-informed neural network for ultrasound nondestructive quantification of  
 664 surface breaking cracks. *Journal of Nondestructive Evaluation*, 39(3):61, 2020.

665

666 Shashank Subramanian, Peter Harrington, Kurt Keutzer, Wahid Bhimji, Dmitriy Morozov, Michael W.  
 667 Mahoney, and Amir Gholami. Towards foundation models for scientific machine learning: Char-  
 668 acterizing scaling and transfer behavior. In *Thirty-seventh Conference on Neural Information  
 669 Processing Systems*, 2023.

670 Makoto Takamoto, Timothy Praditia, Raphael Leiteritz, Daniel MacKinlay, Francesco Alesiani, Dirk  
 671 Pflüger, and Mathias Niepert. Pdebench: An extensive benchmark for scientific machine learning.  
 672 *Advances in Neural Information Processing Systems*, 35:1596–1611, 2022.

673

674 Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh  
 675 Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn  
 676 high frequency functions in low dimensional domains. *Advances in neural information processing  
 677 systems*, 33:7537–7547, 2020.

678 Eleuterio F Toro, Michael Spruce, and William Speares. Restoration of the contact surface in the  
 679 hll-riemann solver. *Shock waves*, 4(1):25–34, 1994.

680

681 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée  
 682 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and  
 683 efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.

684

685 Alasdair Tran, Alexander Mathews, Lexing Xie, and Cheng Soon Ong. Factorized fourier neural  
 686 operators. *arXiv preprint arXiv:2111.13802*, 2021.

687

688 Bram Van Leer. Towards the ultimate conservative difference scheme. v. a second-order sequel to  
 689 godunov’s method. *Journal of computational Physics*, 32(1):101–136, 1979.

690

691 Stephen Walker. On sufficient conditions for bayesian consistency. *Biometrika*, 90(2):482–488, 2003.

692

693 Stephen Walker. New approaches to bayesian consistency. *The Annals of Statistics*, 32(5), October  
 694 2004a. ISSN 0090-5364. doi: 10.1214/009053604000000409.

695

696 Stephen G Walker. Modern bayesian asymptotics. *Statistical Science*, pp. 111–117, 2004b.

697

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

700

701 Hyowon Wi, Jeongwhan Choi, and Noseong Park. Learning advanced self-attention for linear  
 702 transformers in the singular value domain. *arXiv preprint arXiv:2505.08516*, 2025.

703

704 Jared Willard, Xiaowei Jia, Shaoming Xu, Michael Steinbach, and Vipin Kumar. Integrating scientific  
 705 knowledge with machine learning for engineering and environmental systems. *ACM Computing  
 706 Surveys*, 55(4):1–37, 2022.

702 Liu Yang and Stanley J. Osher. Pde generalization of in-context operator networks: A study on 1d  
703 scalar nonlinear conservation laws. *Journal of Computational Physics*, pp. 113379, 2024. ISSN  
704 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2024.113379>.

705

706 Liu Yang, Xuhui Meng, and George Em Karniadakis. B-pinns: Bayesian physics-informed neural  
707 networks for forward and inverse pde problems with noisy data. *Journal of Computational Physics*,  
708 425:109913, 2021.

709

710 Liu Yang, Siting Liu, Tingwei Meng, and Stanley J Osher. In-context operator learning with data  
711 prompts for differential equation problems. *Proceedings of the National Academy of Sciences*, 120  
712 (39):e2310142120, 2023.

713

714 Liu Yang, Siting Liu, and Stanley J Osher. Fine-tune language models as multi-modal differential  
715 equation solvers. *Neural Networks*, pp. 107455, 2025.

716

717 Zhanhong Ye, Xiang Huang, Leheng Chen, Hongsheng Liu, Zidong Wang, and Bin Dong. Pdeformer:  
718 Towards a foundation model for one-dimensional partial differential equations. *arXiv preprint*  
719 *arXiv:2402.12652*, 2024.

720

721 Lei Yuan, Yi-Qing Ni, Xiang-Yun Deng, and Shuo Hao. A-pinn: Auxiliary physics informed neural  
722 networks for forward and inverse problems of nonlinear integro-differential equations. *Journal of*  
723 *Computational Physics*, 462:111260, 2022.

724

725 Ruiqi Zhang, Spencer Frei, and Peter L. Bartlett. Trained transformers learn linear models in-context.  
726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

756 **A THE PROOF OF THEOREM 2.1**  
 757

758 In the following lemma, we first prove for  $D_n$ , i.e., noise-free. We then extend it to the case where  
 759 the noisy data  $\tilde{D}_n$  is used for training.  
 760

761 **Lemma A.1.** *Suppose that for any  $\epsilon > 0$ , there exists a Transformer parameterized by  $\hat{\theta}$  such that*

$$762 \mathbb{E}_c [KL(p_{\hat{\theta}}(\cdot | c, D_n), q_n(\cdot | c, D_n))] < \epsilon, \\ 763$$

764 *for any realization of  $D_n$ . If the posterior consistency condition Eq. (1) holds, and for any  $q \in \mathcal{Q}$ ,  
 765  $p_q(c) = p_\pi(c)$  almost everywhere on  $\mathcal{X}$ , then the following holds almost surely:*

$$766 \mathbb{E}_c [H(p_{\hat{\theta}}(\cdot | c, D_n), p_\pi(\cdot | c))] \xrightarrow{n \rightarrow \infty} 0. \\ 767$$

768 *Proof.* For any two probability distributions  $p$  and  $q$ , recall that the Hellinger distance satisfies  
 769

$$770 H(p, q)^2 = \frac{1}{2} \int (\sqrt{p(\theta)} - \sqrt{q(\theta)})^2 d\theta \leq \frac{1}{2} KL(p \| q). \\ 771$$

772 Thus, if we can show that the Kullback–Leibler divergence  $KL(\hat{p}_{\hat{\theta}}(\cdot | c, \tilde{D}) \| p_\pi(\cdot | c))$  vanishes  
 773 asymptotically, then the convergence in Hellinger distance follows. For any  $n, \epsilon$ , we derive that  
 774

$$775 \mathbb{E}_c [H(p_{\hat{\theta}}(\cdot | c, D_n), p_\pi(\cdot | c))] \stackrel{(1)}{\leq} \mathbb{E}_c [H(p_{\hat{\theta}}(\cdot | c, D_n), q_n(\cdot | c, D_n))] \\ 776 + \mathbb{E}_c [H(q_n(\cdot | c, D_n), p_\pi(\cdot | c))] \\ 777 \stackrel{(2)}{\leq} \sqrt{\frac{1}{2} \mathbb{E}_c [KL(p_{\hat{\theta}}(\cdot | c, D_n), q_n(\cdot | c, D_n))]} \\ 778 + \mathbb{E}_c [H(q_n(\cdot | c, D_n), p_\pi(\cdot | c))] \\ 779 \stackrel{(3)}{\leq} \sqrt{\frac{\epsilon}{2}} + \mathbb{E}_c \left[ 1 - \int_{\mathcal{Y}} \sqrt{\int p_q(y | c) p_\pi(y | c) d\Pi^n(q)} dy \right]^{1/2} \\ 780 \\ 781 \stackrel{(4)}{\leq} \sqrt{\frac{\epsilon}{2}} + \left[ 1 - \int_{\mathcal{C}} \int_{\mathcal{Y}} \sqrt{\int p_q(y, c) p_\pi(y, c) d\Pi^n(q)} dy dc \right]^{1/2} \\ 782 \\ 783 \stackrel{(5)}{=} \sqrt{\frac{\epsilon}{2}} + \left[ \int_{\{q: H(p_\pi, p_q) > \epsilon\}} H(p_q, p_\pi) d\Pi^n(q) \right]^{1/2} \\ 784 \\ 785 \stackrel{(6)}{=} \sqrt{\frac{\epsilon}{2}} + \left[ \int_{\{q: H(p_\pi, p_q) \leq \epsilon\}} H(p_q, p_\pi) d\Pi^n(q) \right]^{1/2} \\ 786 \\ 787 \stackrel{(7)}{=} \sqrt{\frac{\epsilon}{2}} + (\Pi^n(\{q: H(p_\pi, p_q) > \epsilon\}) + \epsilon)^{1/2} \rightarrow \sqrt{\frac{\epsilon}{2}} + \sqrt{\epsilon} \quad \text{a.s.}$$

805  
 806  
 807 The first inequality  $\stackrel{(1)}{\leq}$  is derived from the triangle inequality for the Hellinger distance, which states  
 808 that for any intermediate distribution  $q(\cdot | c, D_n)$ , we have  
 809

$$H(p_{\hat{\theta}}(\cdot | c, D_n), p_\pi(\cdot | c)) \leq H(p_{\hat{\theta}}(\cdot | c, D_n), q_n(\cdot | c, D_n)) + H(q_n(\cdot | c, D_n), p_\pi(\cdot | c)).$$

810  
811 The second inequality  $(\leq)$  uses the fact that the Hellinger distance  $H(p, q)$  is bounded above by the  
812 square root of the KL divergence  $KL(p \parallel q)$ , such that

$$813 \quad 814 \quad H(p, q)^2 \leq \frac{1}{2}KL(p \parallel q). \\ 815$$

816 Thus, we can bound the Hellinger distance by the KL divergence. In the third inequality  $(\leq)$ , we  
817 make use of the assumption

$$818 \quad 819 \quad \mathbb{E}_c [KL(p_{\hat{\theta}}(\cdot \mid c, D_n), q_n(\cdot \mid c, D_n))] < \epsilon,$$

820 and utilize the definition of the Hellinger distance. In  $(\stackrel{(4)}{=})$ , we partition the domain into two regions—  
821 one where the Hellinger distance  $H(p_{\pi}, p_q)$  exceeds  $\epsilon$  and another where it is less than or equal to  
822  $\epsilon$ —and use this partitioning to demonstrate the inequality.

823 Finally, in  $(\stackrel{(5)}{=})$ , by posterior consistency, the region where the Hellinger distance is greater than  $\epsilon$   
824 vanishes as  $n \rightarrow \infty$  such that

$$825 \quad 826 \quad \Pi^n \{q : H(p_{\pi}, p_q) > \epsilon\} \rightarrow 0 \quad \text{almost surely.}$$

827 Since  $\epsilon$  is arbitrary, we can conclude that

$$828 \quad 829 \quad \mathbb{E}_c [H(p_{\hat{\theta}}(\cdot \mid c, D_n), p_{\pi}(\cdot \mid c))] \xrightarrow{n \rightarrow \infty} 0 \quad \text{almost surely.}$$

□

830 Based on Lemma A.1, we can prove Theorem 2.1 as follows.

831 **Theorem A.2.** *Let  $D_n$  be a set of ground-truth prior data, whose size is  $n$ , and  $\tilde{D}_n = D_n + \eta_n$ ,  
832 where  $\eta_n$  is a zero-mean noise distribution with a finite variance, be our observation. Therefore,  $\tilde{D}_n$   
833 is an unbiased observation of  $D_n$ . Let  $p_{\pi}(\cdot \mid c)$  denote the true posterior, and let  $\hat{p}_{\hat{\theta}}(\cdot \mid c, \tilde{D}_n)$  be the  
834 corresponding learned (approximate) posterior for some neural network parameter  $\hat{\theta}$ . Suppose that  
835 the same conditions as in Lemma A.1 hold (with  $\tilde{D}$  in place of  $D_n$ ). Then, for any  $\epsilon > 0$ , it holds that*

$$836 \quad 837 \quad \lim_{n \rightarrow \infty} \mathbb{E}_c [H(\hat{p}_{\hat{\theta}}(\cdot \mid c, \tilde{D}_n), p_{\pi}(\cdot \mid c))] = 0 \quad \text{almost surely,}$$

838 where,  $H(\cdot, \cdot)$  denotes the Hellinger distance.

839 *Proof.* In the noise-free setting of the previous lemma, it is established that

$$840 \quad 841 \quad \mathbb{E}_c [H(\hat{p}_{\hat{\theta}}(\cdot \mid c, D_n), p_{\pi}(\cdot \mid c))] \rightarrow 0 \quad \text{as } n \rightarrow \infty \quad (\text{a.s.}).$$

842 In our case, since  $\tilde{D}_n = D_n + \eta_n$  with  $\mathbb{E}[\eta_n] = 0$ , every occurrence of  $D_n$  is replaced by  $\tilde{D}_n$ , and the  
843 additional noise introduces an extra expectation  $\mathbb{E}_{\eta_n}[\cdot]$  in the relevant integrals. More precisely, one  
844 has

$$845 \quad 846 \quad KL(\hat{p}_{\hat{\theta}}(\cdot \mid c, \tilde{D}) \parallel q_n(\cdot \mid c, \tilde{D})) = KL(\hat{p}_{\hat{\theta}}(\cdot \mid c, D_n + \eta_n) \parallel q_n(\cdot \mid c, D_n + \eta_n)).$$

847 Because  $\eta_n$  is unbiased and of finite variance, the extra terms arising from the noise remain bounded  
848 and can be integrated out without affecting the asymptotic convergence rate established in the noise-  
849 free case. By posterior consistency, the region where  $H(p_{\pi}(\cdot), p_q(\cdot)) > \epsilon$  has vanishing measure  
850 as  $n \rightarrow \infty$ , i.e.  $\Pi \{q : H(p_{\pi}, p_q) > \epsilon\} \rightarrow 0$  almost surely. Uniform integrability and continuity  
851 arguments used to control the posterior concentration remain valid, as the additional terms contributed  
852 by  $\eta_n$  are absorbed by the outer expectation  $\mathbb{E}_{\eta_n}[\cdot]$ . Thus, one concludes that

$$853 \quad 854 \quad \lim_{n \rightarrow \infty} \mathbb{E}_c [H(\hat{p}_{\hat{\theta}}(\cdot \mid c, \tilde{D}), p_{\pi}(\cdot \mid c))] = 0 \quad \text{almost surely.}$$

855 This completes the proof that the consistency result of Lemma A.1 extends to the case where the  
856 prior data is perturbed by small unbiased noise. □

864 **B DATASETS**865 **B.1 FAMILY OF 2D CONVECTION-DIFFUSION-REACTION EQUATIONS**

866 The formulation of the family of 2D CDR equations with three reaction terms, used in Section 4.2, is  
 867 given as follows.

$$870 \text{ 2D CDR: } u_t + \beta \cdot \nabla u - \nabla \cdot (\nu \nabla u) - \sum_{j=1}^3 \rho_j f_j(u) = 0, \quad x \in \mathcal{X}, \quad t \in \mathcal{T},$$

871 where  $\mathcal{X} = [0, \pi] \times [0, \pi]$  and  $\mathcal{T} = [0, 1]$  are solution domain with a resolution of  $32 \times 32 \times 11$ . For the  
 872 family of 2D CDR equations, we use analytic solutions whenever they are available. When an analytic  
 873 solution is not available, the solution is generated using the spectral method, following the same  
 874 approach as in (Cho et al., 2023). The solutions are computed over the time interval  $t \in [0.0, 1.0]$   
 875 with 1001 time steps at intervals of 0.001, from which 21 time steps at intervals of 0.05 (from  $t = 0.0$   
 876 to  $t = 1.0$ ) are used in the experiments. We pre-train our model on the family of 2D CDR equations  
 877 with coefficients  $(\beta_x, \beta_y, \nu_x, \nu_y, \rho_1, \rho_2, \rho_3 \in 0.0, 0.5, 1.0)$ . This results in a total of 2,187 unique  
 878 PDEs. For PINN prior, because a PINN can be trained on only a single PDE at a time, we train an  
 879 individual PINN for each of the 2,187 distinct 2D CDR equations.

883 **B.2 SHALLOW WATER EQUATION**

884 The 2D shallow water equation (SWE) dataset used in Section 4.3.1 is obtained from  
 885 PDEBench (Takamoto et al., 2022)<sup>3</sup>. The SWE is widely used as a benchmark dataset and cor-  
 886 responds to a hyperbolic PDE of the following form:

$$887 \begin{aligned} \partial_t h + \partial_x(hu) + \partial_y(hv) &= 0, \\ 888 \partial_t(hu) + \partial_x \left( u^2 h + \frac{1}{2} g_r h^2 \right) + \partial_y(uvh) &= -g_r h \partial_x b, \\ 889 \partial_t(hv) + \partial_y \left( v^2 h + \frac{1}{2} g_r h^2 \right) + \partial_x(uvh) &= -g_r h \partial_y b, \end{aligned}$$

890 where  $u$  and  $v$  denote the horizontal and vertical velocities,  $h$  is the water depth, and  $g_r$  is the  
 891 gravitational acceleration. The SWE is derived from the general Navier–Stokes equations and is well  
 892 suited for modeling free-surface flow problems.

893 The dataset contains simulation results for the water height  $h$  and is generated using the Py-  
 894 Claw (Ketcheson et al., 2012) Python package. It consists of samples produced from 1,000 different  
 895 initial conditions. The simulation domain is  $\mathcal{X} = [-2.5, 2.5]^2$  over the time interval  $\mathcal{T} = [0.0, 1.0]$ ,  
 896 with a resolution of  $128 \times 128 \times 101$ . From this, we subsample 21 time steps at intervals of 0.05 for  
 897 use in our experiments.

903 **B.3 2D COMPRESSIBLE NAVIER-STOKES EQUATION**

904 The 2D compressible Navier–Stokes equations dataset, referred to as CNSE, used in Section 4.3.2  
 905 was obtained from PDEBench. The Navier–Stokes equation is one of the most widely adopted  
 906 benchmarks for evaluating SciML PDE solvers, and in this work we focus on the compressible fluid  
 907 case. The governing equations are given by

$$908 \begin{aligned} \partial_t \rho + \nabla \cdot (\rho \mathbf{v}) &= 0, \\ 909 \rho (\partial_t \mathbf{v} + \mathbf{v} \cdot \nabla \mathbf{v}) &= -\nabla p + \eta \Delta \mathbf{v} + \left( \zeta + \frac{\eta}{3} \right) \nabla (\nabla \cdot \mathbf{v}), \\ 910 \partial_t \left( \epsilon + \frac{1}{2} \rho v^2 \right) + \nabla \cdot \left[ \left( \epsilon + p + \frac{1}{2} \rho v^2 \right) \mathbf{v} - \mathbf{v} \cdot \boldsymbol{\sigma}' \right] &= 0, \end{aligned}$$

911 where  $\rho$  denotes the mass density,  $\mathbf{v}$  the velocity,  $p$  the gas pressure,  $\epsilon = p/(\Gamma - 1)$  the internal  
 912 energy with  $\Gamma = 5/3$ ,  $\boldsymbol{\sigma}'$  the viscous stress tensor, and  $\eta, \zeta$  the shear and bulk viscosity, respectively.

913 <sup>3</sup>Dataset source: <https://github.com/pdebench/PDEBench>

Numerical solutions were computed in the spatial domain  $\mathcal{X} = [0.0, 1.0]^2$  and temporal domain  $\mathcal{T} = [0.0, 1.0]$  using a second-order HLLC scheme (Toro et al., 1994) with the MUSCL method (Van Leer, 1979) for the inviscid part, and a central difference scheme for the viscous part. We use the solution computed by setting the Mach number as  $M = |\mathbf{v}|/c_s = 1.0$ , where  $c_s = \Gamma p/\rho$  is the sound velocity, and both viscosity coefficients as  $\eta = \zeta = 0.1$ . Periodic boundary conditions were applied, and distinct initial conditions were generated from random fields. The dataset contains simulation results for  $x$ - and  $y$ -velocities  $\mathbf{v}$ , pressure  $p$ , and density  $\rho$ . The dataset consists of 1,000 samples simulated from different initial conditions, originally computed on a  $128 \times 128 \times 21$  grid. Each feature in the dataset was normalized individually for the experiment. For our experiments, we employed a subsampled version of the dataset with resolution  $64 \times 64 \times 5$ . In the experiment we consider a field  $\mathbf{u}(x, t) = [\mathbf{v}(x, t), p(x, t), \rho(x, t)] \in \mathbb{R}^4$  for  $x \in \mathcal{X}$ ,  $t \in \mathcal{T}$ . Here  $\mathbf{u}(x, t)$  denotes a four-dimensional vector field representing four features of the system and the initial condition  $\mathbf{u}(x, 0.0) = \mathbf{u}_0(x)$  for  $x \in \mathcal{X}$ . We are interested in learning the trajectory-predicting operator  $\mathcal{G}^\dagger : \mathbf{u}_0(x) \mapsto (\mathbf{u}(x, 0.25), \mathbf{u}(x, 0.5), \mathbf{u}(x, 0.75), \mathbf{u}(x, 1.0))$ , which maps the initial state  $\mathbf{u}_0 \in L^2((0.0, 1.0)^2; \mathbb{R}^4)$  to the solution at four future time instances.

#### B.4 AIRFOIL

The Airfoil dataset used in Section H.2 was obtained from the dataset provided by (Li et al., 2023)<sup>4</sup>. This dataset was originally generated in (Pfaff et al., 2020) by solving the Euler equations for compressible flow using the finite volume method built into the SU2 library (Palacios et al., 2013)<sup>5</sup>. The Euler equations have the following formulation:

$$\begin{aligned} \partial_t \rho + \nabla \cdot (\rho \mathbf{v}) &= f_1, \\ \partial_t (\rho \mathbf{v}) + \nabla \cdot (\rho \mathbf{v} \otimes \mathbf{v} + p \mathbb{I}) &= \mathbf{f}_2, \\ \partial_t (\rho E) + \nabla \cdot (\rho E \mathbf{v} + p \mathbf{v}) &= f_3, \\ \rho &:= \rho(x, t), \quad \mathbf{v} := \mathbf{v}(x, t), \quad p := p(x, t), \quad x \in \mathcal{X}, \quad t \in [0, T], \end{aligned}$$

where  $\rho$  denotes the density,  $\mathbf{v}$  the velocity field,  $p$  the pressure,  $E$  the total energy per unit mass, and  $f_1, \mathbf{f}_2, f_3$  are generic source terms. Li et al. (2023) formulated this problem as a non-Markovian initial value problem and used a significantly larger time step size. The dataset spans  $t \in \mathcal{T} = [0.0, 4.8]$  with 101 time steps, each spaced by a time interval of 0.24. At each timepoint, the dataset contains values for 5,233 irregular mesh nodes, including node position, node type, velocity, pressure, and density. Similar to CNSE, we consider a field  $\mathbf{u}(x, t) = [\mathbf{v}(x, t), p(x, t), \rho(x, t)] \in \mathbb{R}^4$  for  $x \in \mathcal{X}$ ,  $t \in \mathcal{T}$ . Here  $\mathbf{u}(x, t)$  denotes a four-dimensional vector field representing four features of the system. The initial condition is given by  $\mathbf{u}(x, 0) = \mathbf{u}_0(x)$  for  $x \in \Omega$ . We want to learn the trajectory-predicting operator  $\mathcal{G}^\dagger : \mathbf{u}_0(x) \mapsto (\mathbf{u}(x, 1.2), \mathbf{u}(x, 2.4), \mathbf{u}(x, 3.6), \mathbf{u}(x, 4.8))$ , which maps the initial state  $\mathbf{u}_0 \in L^2(\mathcal{X}; \mathbb{R}^4)$  to the solution at four future time instances.

#### B.5 DARCY FLOW

The Darcy Flow dataset used in Section H.3 is also obtained from PDEBench. Darcy Flow is a widely adopted benchmark for evaluating operator learning models, formulated as a time-independent elliptic PDE. The governing equation is given by:

$$\begin{aligned} -\nabla(a(x)\nabla u(x)) &= f(x), \quad x \in \mathcal{X}, \\ u(x) &= 0, \quad x \in \partial\mathcal{X}, \end{aligned}$$

where the source term  $f(x)$  is fixed to a constant value of 0.1. The dataset consists of steady-state solutions  $u$  corresponding to the viscosity fields  $a$ . To obtain these solutions, we simulate the following time-dependent formulation with random initial conditions until convergence to steady state:

$$\partial_t u(x, t) - \nabla(a(x)\nabla u(x, t)) = f(x), \quad x \in \mathcal{X}$$

Numerical solutions are simulated using a second-order central difference scheme in both time and space. Each data sample is represented as an input–output pair, where the input is a viscosity field

<sup>4</sup>Dataset Source: <https://github.com/BaratiLab/OFormer>

<sup>5</sup>Dataset and numerical solver source: <https://github.com/merantix-momentum/gnn-bvp-solver>

$a(x)$  and the output is the corresponding steady-state solution  $u(x)$ . Both inputs and outputs are discretized on a  $128 \times 128$  spatial grid. The dataset contains 2,000 samples, each generated from distinct random initial conditions within the simulation domain  $\mathcal{X} = [0.0, 1.0]^2$ . Unlike the CNSE setting, here we aim to learn the operator mapping a viscosity field  $a(x) \in L^2((0.0, 1.0)^2; \mathbb{R})$  to the corresponding steady-state solution  $u(x)$  for  $x \in \mathcal{X}$ . Specifically, we want to learn the operator  $\mathcal{G}^\dagger : a(x) \mapsto u(x)$ .

The specific information about PDE types used in each experiment is described in Table 6. Moreover, the detailed values regarding how many training, validation, and testing data points in each experiment are provided in Appendix I. For the spatiotemporal interpolation with unseen coefficients in the family of 2D CDR equation experiments, we use the same time range as in the corresponding task with seen coefficients.

Table 6: Time range of data points used for training, validation, and testing in all experiments.

PDE (Section)	Task	Train	Valid		Test	
			Context	Query	Context	Query
Family of 2D CDR equations (Section 4.2)	Spatiotemporal interpolation	$t \in [0.0, 0.5]$	predefined in $t \in [0.0, 0.5]$	grid at $t = 0.25$	predefined in $t \in [0.0, 0.5]$	grids at $t \in \{0.05, 0.15, 0.35, 0.45\}$
	Temporal extrapolation	$t \in [0.0, 0.5]$	grid at $t = 0.5$	grid at $t = 0.6$	grid at $t = 0.6$	grids at $t \in [0.6, 1.0]$
	Unseen coeff temporal extrapolation	-	-	-	grid at $t = 0.1$	grids at $t \in [0.2, 0.5]$
SWE (Section 4.3.1)	Spatiotemporal interpolation	$t \in [0.0, 0.5]$	predefined in $t \in [0.0, 0.5]$	grid at $t = 0.25$	predefined in $t \in [0.0, 0.5]$	grids at $t \in \{0.05, 0.15, 0.35, 0.45\}$
	Temporal extrapolation	$t \in [0.0, 0.5]$	grid at $t = 0.5$	grid at $t = 0.6$	grid at $t = 0.6$	grids at $t \in [0.6, 1.0]$
CNSE (Section 4.3.2)	Operator learning	$t \in \{0.0, 0.25, 0.5, 0.75, 1.0\}$	grid at $t = 0.0$	grids at $t \in \{0.25, 0.5, 0.75, 1.0\}$	grid at $t = 0.0$	grids at $t \in \{0.25, 0.5, 0.75, 1.0\}$
Airfoil (Section H.2)	Operator learning	$t \in \{0.0, 1.2, 2.4, 3.6, 4.8\}$	grid at $t = 0.0$	grids at $t \in \{1.2, 2.4, 3.6, 4.8\}$	grid at $t = 0.0$	grids at $t \in \{1.2, 2.4, 3.6, 4.8\}$
Darcy Flow (Section H.3)	Operator learning			$a \mapsto u$		

## C PINN USED IN PRIOR GENERATION

In this study, we utilize the PINN introduced by Raissi et al. (2019a) to generate PINN priors. The loss function employed during the training of the PINN is as follows:

$$\mathcal{L} = \mathcal{L}_u + \mathcal{L}_f + \mathcal{L}_b,$$

where  $\mathcal{L}_u$ ,  $\mathcal{L}_f$  and  $\mathcal{L}_b$  are defined as

$$\mathcal{L}_u = \frac{1}{N_u} \sum (\tilde{u}(x, 0) - u(x, 0))^2, \quad \mathcal{L}_f = \frac{1}{N_f} \sum (\mathcal{N}(t, x, u, \alpha))^2, \quad \mathcal{L}_b = \frac{1}{N_b} \sum (\tilde{u}(0, t) - \tilde{u}(2\pi, t))^2,$$

for  $N_u$  points at initial condition,  $N_f$  collocation points, and  $N_b$  boundary points.

The generation time for the PINN prior varies depending on the number of training epochs and the PINN loss threshold. Based on the configuration used in our experiments, it took approximately 373 seconds to generate the prior for a single coefficient combination of the 2D CDR and utilized up to 4,552 MB of memory.

## D DETAILED DESCRIPTION OF THE BASELINES

We provided descriptions of each baseline discussed in Section 4.1.

- DeepONet is a neural operator architecture designed to learn operators mapping input functions to output functions. It combines branch and trunk networks to predict values in a function space.
- FNO learns solution operators for partial differential equations (PDEs) using the Fourier transform. By mapping inputs to a frequency domain, FNO captures complex patterns and long-range dependencies and models complex systems.

- 1026 • F-FNO extends the Fourier Neural Operator by factorizing its layers to reduce computational  
1027 costs. This factorization enables efficient learning of solution operators for complex systems.
- 1028 • A-FNO is a variant of the Fourier Neural Operator that dynamically adjusts the resolution  
1029 of the frequency domain during training. This adaptation aims to capture relevant features  
1030 across scales, enabling more flexible modeling of complex systems.
- 1031 • Poseidon is a multiscale operator transformer architecture designed to model solution op-  
1032 erators for PDEs with continuous-in-time capability. It introduces time-conditioned layer  
1033 normalization and processes multi-resolution input features to capture complex spatiotem-  
1034 poral dynamics.
- 1035 • DPOT is a scalable neural operator model designed for PDEs, featuring a Fourier attention  
1036 mechanism that enables efficient learning of complex spatiotemporal dynamics. Its architec-  
1037 ture incorporates an auto-regressive denoising pre-training strategy, facilitating stable and  
1038 flexible modeling across diverse PDE datasets with varying resolutions and geometries.

1040 In the spatiotemporal interpolation task, models other than DeepONet are excluded from comparison  
1041 because they cannot process mesh-structured inputs. DeepONet requires input data from fixed  
1042 positions for each dataset; therefore, for a fair comparison, our model is also provided with context  
1043 data from the same positions during evaluation. For Poseidon and DPOT, which are scientific  
1044 foundation models providing pre-trained weights, we utilized their released weights and performed  
1045 fine-tuning. For a fair comparison, we trained them and our model using only half the number of  
1046 epochs compared to the other models. Furthermore, computational costs, memory usage, experiment  
1047 settings, and hyperparameters are comprehensively outlined in Appendix I.

1048 Table 7: Major comparisons between baselines and our model on the shape of input and target  
1049 data (context and query in our model). Neural operator baselines (excluding DeepONet) can only  
1050 predict targets at predetermined locations and are limited to processing grid-structured input data.  
1051 DeepONet, on the other hand, can handle inputs of various shapes but requires fixed-coordinate input  
1052 values, though it can produce target solutions at desired coordinates. In contrast, our model imposes  
1053 no restrictions on the shape, coordinates, or number of inputs and can generate target solutions at  
1054 arbitrary coordinates.

Data	FNO & F-FNO & A-FNO & Poseidon & DPOT	DeepONet	Ours
Input data	Grid of predefined coordinate	Mesh of predefined coordinate	<b>Mesh of arbitrary coordinate</b>
Target data	Grid of predefined coordinate	<b>Mesh of arbitrary coordinate</b>	<b>Mesh of arbitrary coordinate</b>

## E MODEL ARCHITECTURE

1063 Our model is fundamentally based on the Prior-Fitted Network (PFN) architecture (Müller et al.,  
1064 2022), and consists of three main components: an encoder, a Transformer block, and a decoder. In  
1065 this section, we describe the encoder. For clarity, we denote the training contexts and queries as  $D$   
1066 and  $T$ , respectively, consistent with the main text. Each input in  $D$  and  $T$  is composed of spatial  
1067 coordinates  $x$ , temporal coordinates  $t$ , and solution values  $y$ ; for queries  $T$ , the solution values are  
1068 masked. The encoder enriches these inputs through a Fourier feature embedding (FFE) followed by a  
1069 multilayer perceptron (MLP).

1070 **Encoder** The role of FFE is to augment the raw inputs with high-frequency components, enabling  
1071 the model to better capture complex solution patterns. This advantage of using FFE in conjunction  
1072 with MLPs has been verified in (Tancik et al., 2020; Ma et al., 2025). For grid-based inputs, we can  
1073 employ a discrete Fourier transform via fast Fourier transform (FFT). For coordinate-based inputs, we  
1074 use sinusoidal feature mappings with predetermined frequencies. Specifically, given an input  $v \in \mathbb{R}$ ,  
1075 we define the embedding as

$$1076 \text{FFE}(v; n) = [\sin(v \cdot \omega_1), \cos(v \cdot \omega_1), \dots, \sin(v \cdot \omega_n), \cos(v \cdot \omega_n)],$$

1077 where the frequency coefficients are set as  $\omega_i = 2\pi/i$ . After applying FFE, the encoder constructs  
1078 three types of feature vectors:  $l_{CD} := [x_D, t_D, \text{FFE}(x_D), \text{FFE}(t_D)]$  for context domain information,

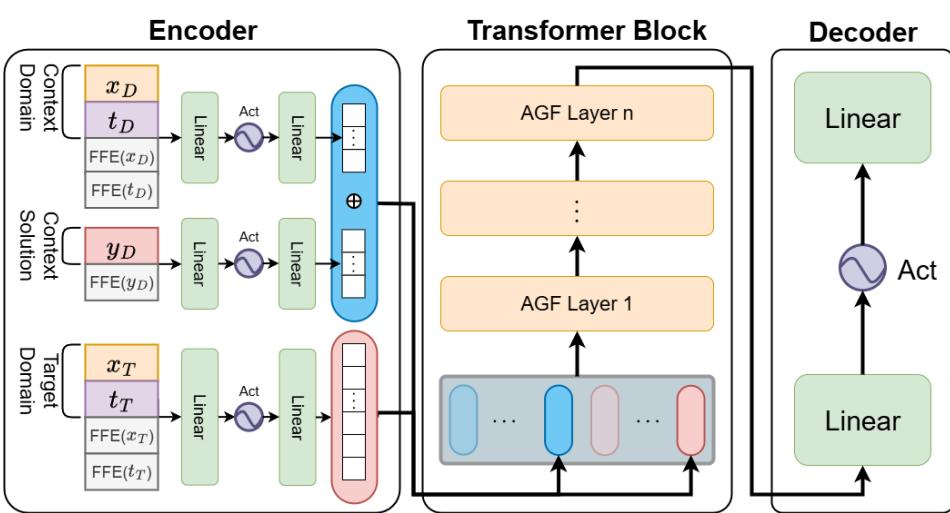


Figure 3: A diagram of PDE-PFN’s architecture. The diagram is presented as an extension of Figure 2. The model consists of three main components: an encoder, a Transformer block, and a decoder. FFE denotes the Fourier feature embedding, and Act denotes the activation function.

$l_{CS} := [y_D, \text{FFE}(y_D)]$  for context solutions, and  $l_{QD} := [x_T, t_T, \text{FFE}(x_T), \text{FFE}(t_T)]$  for query domain information, by concatenating the raw inputs and FFE results. Each of these vectors is passed through a separate MLP to learn PDE-specific embeddings. Each MLP consists of two linear layers and a single activation function between linear layers. For the activation, we employ a rational activation function (Molina et al., 2019), which has an adaptive nature and has been shown to provide greater flexibility compared to conventional nonlinearities. Finally,  $l_{CD}$  and  $l_{CS}$  are concatenated after passing each MLP to form the context representation used by the Transformer block.

**Transformer block** The encoder output is then processed through a Transformer block, which consists of a sequence of Transformer layers. Instead of the *vanilla* Transformer layer, we adopt the attentive graph filter (AGF) layer Wi et al. (2025), which provides greater flexibility while reducing computational complexity. The AGF layer enhances standard linear attention by learning attention in the singular value domain. Specifically, it decomposes the attention operation into a set of spectral components and adaptively learns their importance, enabling the model to capture the dependencies more effectively. Specifically, the attention matrix  $A \in \mathbb{R}^{n \times m}$  can be decomposed using singular value decomposition (SVD) as  $A = U\Sigma V^T$ , where  $U$  and  $V$  are orthogonal matrices, and  $\Sigma = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_r)$  contains the singular values. Therefore, as discussed in Section 3, when using the AGF layer, we include a regularization term to ensure the orthogonality of the matrices  $U$  and  $V$ . The regularization term is defined as follows:

$$\mathcal{L}_{\text{AGF}} = \frac{1}{n^2} (\|U^\top U - \mathbf{I}\|) + \frac{1}{m^2} (\|V V^\top - \mathbf{I}\|),$$

where  $\mathbf{I}$  is identity matrix. AGF learns to adaptively reweight these spectral components, effectively filtering the most informative modes while suppressing noise or redundant information. By learning attention in the singular value domain, AGF provides a more powerful representation of global dependencies compared to conventional linear attention, while avoiding the quadratic complexity of the *vanilla* Transformers.

To further enrich the model capacity, we separate the parameters used for self-attention and cross-attention within each layer. As illustrated in Figure 2, self-attention is applied among the contexts  $D$ , while cross-attention allows queries  $T$  to attend to contexts  $D$ . In the figure, self-attention and cross-attention are represented by red and blue arrows, respectively. The numbers of blue rods  $D$  and red rods  $T$  are described in Table 20 in Appendix I.

**Decoder** The output of the Transformer block is finally passed through a decoder, implemented as a simple MLP. Similar to the encoder, this MLP consists of two linear layers and a single

1134 activation function between linear layers. The decoder maps the latent representations produced by  
 1135 the Transformer into the predicted solution values at the queried coordinates.  
 1136

## 1137 F TRAINING

1139 Unlike other task-specific models, our model is trained in an integrated manner, enabling it to solve  
 1140 multiple tasks simultaneously in that dataset. For cases where both spatiotemporal interpolation and  
 1141 temporal extrapolation tasks must be evaluated, such as in the family of 2D CDR equations and SWE  
 1142 datasets, we perform training with temporal extrapolation as the primary objective. In this setup,  
 1143 other baselines treat grid data points at  $t_n$  and  $t_{n+1}$  as fixed input–output pairs. In contrast, our model  
 1144 constructs pairs by randomly mixing the data points within each  $(t_n, t_{n+1})$  pair independently, rather  
 1145 than preserving their exact correspondence. This strategy enables the model to learn in a way that  
 1146 supports both spatiotemporal interpolation and temporal extrapolation within a unified framework.  
 1147 For the CNSE, Airfoil, and Darcy Flow datasets, where we conducted experiments only on operator  
 1148 learning, random mixing was not applied.  
 1149

1150 **Training algorithm** We train our model as follows:

---

### 1151 Algorithm 1 Training our model

---

1153 1: **Input:** contexts  $D$  and queries  $T$  from prior  $p(\mathcal{D})$  in dataset  
 1154 2: **Output:** Our model  $\tilde{u}_\theta$  which can approximate the PPD  
 1155 3: Initialize the model  $\tilde{u}_\theta$   
 1156 4: **for**  $i = 1$  to  $n$  **do**  
 1157 5:   Sample  $\alpha \in \Omega$  and  $D, T \subseteq \tilde{u}(\alpha) \sim p(\mathcal{D})$  ( $D := \{(x_D^{(i)}, t_D^{(i)})\}_{i=1}^{N_D}$ ,  $T := \{(x_T^{(j)}, t_T^{(j)})\}_{j=1}^{N_T}$ )  
 1158 6:   **if** dataset  $\in \{\text{Family of 2D CDR, SWE}\}$  **then**  
 1159 7:     Regenerate  $D'$  and  $T'$  by shuffling data points in  $D \cup T$  independently within each  $(t_n, t_{n+1})$   
 1160 8:     pair, while preserving  $|D'| = N_D$  and  $|T'| = N_T$   
 1161 9:     **end if**  
 1162 10:   Compute MSE loss  $\mathcal{L}_\alpha = \frac{1}{N_T} \sum_{j=1}^{N_T} \left\{ \tilde{u}_\theta(x_T^{(j)}, t_T^{(j)} | D'_n) - \tilde{u}(x_T^{(j)}, t_T^{(j)}) \right\}^2$   
 1163 11:   Compute AGF regularization term  $\mathcal{L}_{AGF}$  and objective function  $\mathcal{L} = \mathcal{L}_\alpha + \mathcal{L}_{AGF}$   
 1164 12:   Update parameters  $\theta$  with an AdamW optimizer  
 1165 **end for**  
 1166

---

## 1167 G EVALUATION

1169 We employ  $L_1$  absolute,  $L_2$  relative, and  $L_\infty$  relative errors between the model’s prediction for test  
 1170 queries and the ground truth. The errors are then averaged over the target parameter space or test  
 1171 dataset. Given the true solution  $y_{\alpha,k}$  and the corresponding prediction  $\hat{y}_{\alpha,k}$  at the  $k$ -th time point out  
 1172 of a total of  $K$  evaluation time points, each metric is computed as follows:  
 1173

$$L_p \text{ abs error} = \frac{1}{|\Omega| \cdot K} \sum_{\alpha \in \Omega} \sum_{k=1}^K \|y_{\alpha,k} - \hat{y}_{\alpha,k}\|_p, \quad L_p \text{ rel error} = \frac{1}{|\Omega| \cdot K} \sum_{\alpha \in \Omega} \sum_{k=1}^K \frac{\|y_{\alpha,k} - \hat{y}_{\alpha,k}\|_p}{\|y_{\alpha,k}\|_p}.$$

## 1178 H ADDITIONAL EXPERIMENTS

### 1180 H.1 TEST TIME EVALUATION GIVEN PINN PRIOR IN THE FAMILY OF 2D CDR EQUATIONS

1181 As an additional experiment on the family of 2D CDR equations, we modify the test procedure from  
 1182 Section 4.2.1. While keeping all other experimental settings the same, we compare the evaluation  
 1183 results when noisy PINN priors are provided as input during testing. This experiment is designed to  
 1184 assess the robustness of the models in producing accurate solutions despite noisy inputs.  
 1185

1186 As shown in the experimental results presented in Table 8, Ours achieves the best performance,  
 1187 demonstrating strong robustness. For Ours (PINN), however, the performance in terms of two relative  
 1188 errors on the spatiotemporal interpolation task is inferior to that of DeepONet. This can be attributed  
 1189

1188  
 1189  
 1190  
 1191  
 1192 Table 8: The evaluation results for the spatiotemporal interpolation and temporal extrapolation task  
 1193 applied to the family of 2D CDR equations given a noisy PINN prior at evaluation. They are measured  
 1194 at the seen coefficients ( $\beta_x, \beta_y, \nu_x, \nu_y, \rho_1, \rho_2, \rho_3 \in \{0.0, 0.5, 1.0\}$ ). The best performance is marked  
 1195 in **bold** and the second-best performance is marked with underline.  
 1196

Task	Metric	Ours	Ours (PINN)	DeepONet	FNO	F-FNO	A-FNO	<u>LNO</u>	Poseidon	DPOT
Spatiotemporal interpolation	$L_1$ Abs	<b>0.01953</b>	<u>0.02445</u>	0.02540	-	-	-	-	-	-
	$L_2$ Rel	<b>0.02755</b>	<u>0.03467</u>	<u>0.03413</u>	-	-	-	-	-	-
	$L_\infty$ Rel	0.10158	0.10716	<b>0.09286</b>	-	-	-	-	-	-
Temporal extrapolation	$L_1$ Abs	<b>0.01655</b>	<u>0.02099</u>	0.06888	0.03458	0.02816	0.05794	<b>0.03048</b>	0.14018	0.05051
	$L_2$ Rel	<b>0.02486</b>	<u>0.02893</u>	0.08591	0.04268	0.03582	0.07903	<b>0.03907</b>	0.15380	0.08624
	$L_\infty$ Rel	0.08830	<b>0.09333</b>	0.23386	0.10203	0.10677	0.28805	<b>0.11730</b>	0.28481	0.08624

1197  
 1198 to the fact that Ours (PINN) is pre-trained to predict solutions from PINN priors as inputs, which may  
 1199 cause the model to generate outputs closer to the PINN prior rather than the analytic solution when  
 1200 the prior is provided at test time. Nevertheless, since Ours (PINN) still achieves better performance  
 1201 in terms of  $L_1$  absolute error, we can conclude that it also exhibits a certain degree of robustness.  
 1202

## 1203 H.2 OPERATOR LEARNING IN AIRFOIL

1204 To evaluate input flexibility in operator learning, we conduct experiments on the Airfoil dataset  
 1205 defined on an irregular mesh structure. This dataset consists of 2,000 samples, of which 700 are used  
 1206 for training, 100 for validation, and 200 for testing. The target operator is formulated similarly to that  
 1207 in Section 4.3.2: given the initial condition, the model is required to predict the solutions at four time  
 1208 steps,  $t = 1.2, 2.4, 3.6$ , and  $4.8$  (see Appendix B.4 for a formal description). Most of the baselines  
 1209 used in Section 4 rely on data defined on a regular grid, so that cannot be applied here, since the  
 1210 Airfoil dataset is defined on an irregular mesh. Consequently, among the existing baselines, we only  
 1211 use DeepONet and additionally include Oformer (Li et al., 2023) as a new baseline.  
 1212

1213  
 1214 Table 9: The evaluation results for the operator learning task applied to the Airfoil dataset. The best  
 1215 performance is marked in **bold** and the second-best performance is marked with underline.  
 1216

Metric	Ours	Ours (PINN)	DeepONet	Oformer
$L_1$ Abs	0.11047	<b>0.10755</b>	0.23160	0.12039
$L_2$ Rel	<u>0.18831</u>	<b>0.18357</b>	0.44101	0.19744
$L_\infty$ Rel	0.33584	<u>0.32887</u>	0.73076	<b>0.32318</b>

1217 The difference from the experimental results reported in the original Oformer paper arises because  
 1218 their setting involved using the first four time steps as input to predict subsequent time steps,  
 1219 whereas our experiment is formulated based on the initial condition only. The experimental results  
 1220 in Table 9 show that both versions of our model outperform the baselines on two metrics. These  
 1221 results demonstrate that our model’s ICL capability effectively extends to irregular mesh data and can  
 1222 generalize well to new PDE problems. Detailed information on the best hyperparameter settings can  
 1223 be found in Appendix I.3.  
 1224

## 1225 H.3 OPERATOR LEARNING IN DARCY FLOW

1226 To further evaluate the PDE generalization ability of our model in operator learning, we conducted  
 1227 additional experiments on the Darcy Flow dataset. This dataset consists of 2,000 samples, with 1,400  
 1228 used for training, 200 for validation, and 400 for testing. The operator learning task is defined as  
 1229 predicting the steady-state solution  $u(x)$  from a given viscosity field  $a(x)$  (see Appendix B.5 for  
 1230 a formal description). Unlike the CNSE dataset, where both context and query represent the same  
 1231 feature, the Darcy Flow dataset involves mapping between different quantities. This makes it a distinct  
 1232 form of operator learning task.  
 1233

1234  
 1235 Table 10: The evaluation results for the operator learning task applied to Darcy Flow. The best  
 1236 performance is marked in **bold** and the second-best performance is marked with underline.  
 1237

Metric	Ours	Ours (PINN)	DeepONet	FNO	F-FNO	A-FNO	Poseidon	DPOT
$L_1$ Abs	0.00796	<u>0.00742</u>	0.00996	0.00792	0.00806	0.00906	0.00998	<b>0.00726</b>
$L_2$ Rel	0.28771	<b>0.28091</b>	0.40759	0.29292	0.34149	0.40357	0.41001	0.51505
$L_\infty$ Rel	<b>0.39694</b>	0.40124	0.59057	0.46398	0.59381	0.80674	0.52933	0.51505

The experimental results in Table 10 show that both versions of our model outperform the baselines on two metrics. These findings further confirm that our model maintains strong PDE generalization ability even when the operator learning task requires transferring between different physical quantities. Detailed information on the best hyperparameter settings can be found in Appendix I.3.

#### H.4 SENSITIVITY ANALYSIS TO THE NUMBER OF TEST CONTEXT POINTS IN THE FAMILY OF 2D CDR EQUATION

In this section, we provide empirical validation of Theorem 2.1, which establishes the theoretical consistency of the neural network’s posterior predictive distribution (PPD) as the size of the data  $D_n$  increases. Specifically, we evaluate the sensitivity of the neural network’s PPD to the number of the context set  $\bar{D}$  provided during the test process.

To this end, we conduct an evaluation of spatiotemporal interpolation on the family of 2D CDR equations experiment by varying the number of the context set  $\bar{D}$  while keeping other factors unchanged, including the best hyperparameter settings. The results, depicted in Figure 4, clearly demonstrate that as the size of  $\bar{D}$  increases, the error consistently decreases across all three evaluation metrics. This behavior aligns perfectly with the theoretical prediction in Theorem 2.1, where the posterior approximation is shown to converge toward the true distribution as the data size grows.

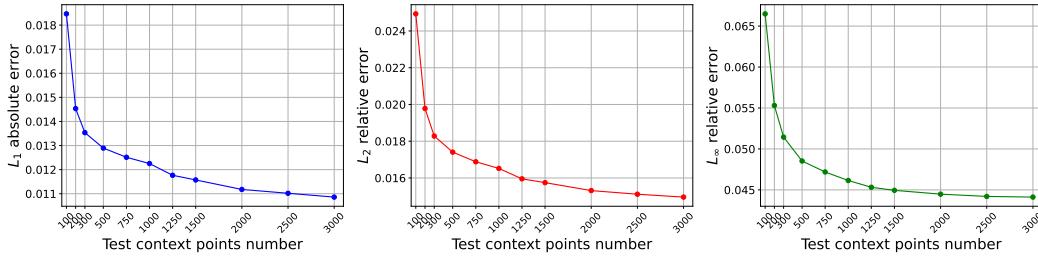


Figure 4: Sensitivity to the number of the context data  $\bar{D}$  during the test.

These experimental findings not only validate the theoretical insights of Theorem 2.1 but also emphasize the robustness and accuracy of the neural network’s PPD approximation under the given prior. The decreasing error trend highlights how the model effectively integrates increasing amounts of data to produce predictions that are more consistent with the true underlying posterior distribution. This synergy between theory and empirical observation strongly supports the reliability and effectiveness of the proposed approach.

#### H.5 UNSEEN DOMAIN TEMPORAL EXTRAPOLATION FOR UNSEEN COEFFICIENT 2D CDR EQUATIONS

In the experiments on unseen coefficients presented in the main text, the spatiotemporal domain was restricted to the temporal range used during training. This design choice ensured that only the parameter domain included unseen values, preventing both parameter and temporal extrapolation from occurring simultaneously. In this section, we extend the evaluation to a more challenging setting in which both the parameter domain and the temporal domain involve unseen ranges. The unseen coefficient ranges follow those used in Section 4.2.2, and for the temporal domain we adopt the evaluation range from the temporal extrapolation task in Section 4.2.1.

Table 11: The evaluation results for the temporal extrapolation tasks applied to the family of 2D CDR equations. They are measured at the unseen coefficients and unseen temporal domain. The best performance is marked in **bold** and the second-best performance is marked with underline.

Coefficients	Task	Metric	Ours	Ours (PINN)	DeepONet	FNO	F-FNO	A-FNO	LNO	Poseidon	DPOT
Interpolation	Temporal extrapolation	$L_1$ Abs	<b>0.00736</b>	<u>0.00939</u>	0.05139	0.10863	0.01405	0.04695	0.01839	0.13852	0.04610
		$L_2$ Rel	<b>0.00954</b>	<u>0.01178</u>	0.06179	0.14380	0.01636	0.06227	0.02049	0.14993	0.04676
		$L_\infty$ Rel	<b>0.03193</b>	<u>0.03693</u>	0.16008	0.34858	0.03847	0.21750	0.04512	0.26293	0.04676
Extrapolation	Temporal extrapolation	$L_1$ Abs	<b>0.01027</b>	<u>0.02063</u>	0.05514	0.03887	0.02312	0.05013	0.02560	0.14661	0.05186
		$L_2$ Rel	<b>0.01068</b>	<u>0.02169</u>	0.06450	0.05114	0.02470	0.06351	0.02677	0.15650	0.05189
		$L_\infty$ Rel	<b>0.01661</b>	<u>0.03474</u>	0.16555	0.13713	0.04591	0.20955	0.04911	0.27502	0.05189

1296 The experimental results in Table 11 show that our model continues to achieve the best performance.  
 1297 In contrast, the baselines experience greater difficulty when both types of extrapolation are required  
 1298 simultaneously, performing noticeably worse than our model in this setting.  
 1299

## 1300 H.6 PDE GENERALIZATION FOR 3D DIFFUSION-REACTION EQUATIONS 1301

1302 To further explore the potential of our model for 3D PDEs, we conduct additional experiments on 3D  
 1303 diffusion–reaction equations. We construct a 3D diffusion–reaction equation dataset by extending  
 1304 the same procedure used for generating the 2D CDR equations. This formulation includes three  
 1305 diffusion coefficients ( $\nu_x, \nu_y, \nu_z$ ) and one reaction coefficient  $\rho$ , each sampled from 0.0, 0.5, 1.0, 1.5,  
 1306 resulting in 256 distinct equations. We perform experiments on the same two tasks as in the 2D CDR  
 1307 equations—spatiotemporal interpolation and temporal extrapolation—using DeepONet, FNO, and  
 1308 DPOT as baselines, since only they can handle 3D datasets among baselines.  
 1309

1310 Table 12: The evaluation results for the spatiotemporal interpolation and temporal extrapolation  
 1311 tasks applied to the 3D diffusion reaction equations. They are measured at the seen coefficients  
 1312 ( $\nu_x, \nu_y, \nu_z, \rho_1 \in \{0.0, 0.5, 1.0, 1.5\}$ ). The best performance is marked in **bold** and the second-best  
 1313 performance is marked with underline.

Task	Metric	Ours	Ours (PINN)	DeepONet	FNO	DPOT
Spatiotemporal interpolation	$L_1$ Abs	<b>0.02685</b>	0.02843	<u>0.02816</u>	-	-
	$L_2$ Rel	<b>0.03711</b>	0.04144	<u>0.03981</u>	-	-
	$L_\infty$ Rel	<b>0.10938</b>	0.12823	<u>0.12724</u>	-	-
Temporal extrapolation	$L_1$ Abs	<b>0.01031</b>	0.01039	0.02421	0.01521	<u>0.01485</u>
	$L_2$ Rel	<b>0.01376</b>	<u>0.01391</u>	0.03033	0.01824	0.01638
	$L_\infty$ Rel	<u>0.04180</u>	0.04545	0.11530	0.05545	<b>0.01638</b>

1314 The experimental results in Table 12 show that our model outperforms the baselines in spatiotemporal  
 1315 interpolation in all metrics and in temporal extrapolation on two metrics, demonstrating its potential  
 1316 to generalize effectively to 3D datasets.  
 1317

## 1318 H.7 ABLATION STUDY 1319

1320 In this section, we conduct an ablation study to analyze the necessity of the key components in  
 1321 our model. Using the 2D CDR equation experiments, we replace AGF with a *vanilla* Transformer  
 1322 layer and substitute the rational activation with GELU to evaluate the impact of each component on  
 1323 performance.  
 1324

1325 Table 13: The evaluation results for the spatiotemporal interpolation and temporal extrapolation  
 1326 tasks applied to the family of 2D CDR equations. “Ours w.o. agf,” “Ours w.o. act,” and “Ours  
 1327 w.o. all” denote replacing AGF with a vanilla Transformer block, replacing the rational activation  
 1328 with GELU, and replacing both, respectively. They are measured at the seen coefficients  
 1329 ( $\beta_x, \beta_y, \nu_x, \nu_y, \rho_1, \rho_2, \rho_3 \in \{0.0, 0.5, 1.0\}$ ). The best performance is marked in **bold** and the second-  
 1330 best performance is marked with underline.  
 1331

Task	Metric	Ours	Ours w.o. agf	Ours w.o. act	Ours w.o. all
Spatiotemporal interpolation	$L_1$ Abs	<b>0.01218</b>	0.01242	<b>0.01040</b>	0.01780
	$L_2$ Rel	0.01643	<u>0.01625</u>	<b>0.01495</b>	0.02281
	$L_\infty$ Rel	0.04600	<b>0.04297</b>	0.04399	0.05736
Temporal extrapolation	$L_1$ Abs	<b>0.01474</b>	0.01776	<u>0.01725</u>	0.02454
	$L_2$ Rel	<b>0.02261</b>	0.02779	<u>0.02665</u>	0.03437
	$L_\infty$ Rel	<b>0.08444</b>	0.09554	<u>0.09376</u>	0.10625

1342 The experimental results in Table 13 show that although the ablated variants may occasionally perform  
 1343 slightly better on the interpolation task, our full model consistently achieves superior performance  
 1344 on the more challenging temporal extrapolation task. This supports our claim that the proposed  
 1345 architectural components enhance the model’s ability to generalize across diverse tasks in a stable  
 1346 and efficient manner.  
 1347

1348 We additionally examine whether the robustness observed in Section 4.2 is preserved by conducting  
 1349 the same PINN-prior evaluation on the “Ours with sinusoidal activation” and “Ours w.o. AGF”  
 1350 variants.

Table 14: The evaluation results for the spatiotemporal interpolation and temporal extrapolation tasks applied to the family of 2D CDR equations. “Ours w.o. agf,” “Ours w.o. act,” and “Ours w.o. all” denote replacing AGF with a vanilla Transformer block, replacing the rational activation with GELU, and replacing both, respectively. They are measured at the seen coefficients ( $\beta_x, \beta_y, \nu_x, \nu_y, \rho_1, \rho_2, \rho_3 \in \{0.0, 0.5, 1.0\}$ ). The best performance is marked in **bold** and the second-best performance is marked with underline.

Task	Metric	Ours	Ours w.o. agf	Ours with sinusoidal
Spatiotemporal interpolation	$L_1$ Abs	<b>0.01507</b>	0.01667	0.01482
	$L_2$ Rel	<b>0.02057</b>	0.02171	<u>0.02096</u>
	$L_\infty$ Rel	<u>0.06584</u>	<b>0.05536</b>	0.06946
Temporal extrapolation	$L_1$ Abs	<b>0.01940</b>	<u>0.02105</u>	0.02671
	$L_2$ Rel	<b>0.02705</b>	<u>0.03030</u>	0.04159
	$L_\infty$ Rel	<b>0.08802</b>	<u>0.09518</u>	0.14922

Table 14 shows that our model achieves the best results on both tasks, confirming strong robustness, whereas both ablations exhibit degraded performance. These results further indicate that the AGF layer and rational activation function contribute significantly to the model’s robustness.

## H.8 TEST TIME EVALUATION GIVEN PRIOR WITH BIASED NOISE IN THE FAMILY OF 2D CDR EQUATIONS

Table 15: The evaluation results for the temporal extrapolation tasks applied to the family of 2D CDR equations. They are measured at the seen coefficients ( $\beta_x, \beta_y, \nu_x, \nu_y, \rho_1, \rho_2, \rho_3 \in \{0.0, 0.5, 1.0\}$ ).

Bias	No noise	0.1%	0.5%	1.0%	1.5%	2.0%	2.5%	3.0%	3.5%	4.0%
$L_1$ Abs	0.01474	0.01483	0.01472	0.01481	0.01514	0.01572	0.01664	0.01800	0.01985	0.02212
$L_2$ Rel	0.02261	0.02271	0.02246	0.02244	0.02271	0.02330	0.02255	0.02573	0.02772	0.03009
$L_\infty$ Rel	0.08444	0.08556	0.08459	0.08426	0.08490	0.08646	0.08901	0.09265	0.09733	0.10276

Since Theorem 2.1 assumes unbiased noise, we further investigate how robust our model remains when the inference stage includes biased Gaussian noise added to the analytic solution. We set the standard deviation of Gaussian noise to  $\sigma = 0.01$  and vary only the magnitude of the bias. The evaluation is performed on the temporal extrapolation task of the 2D CDR equations, using the  $L_2$  relative error for comparison.

The experimental results in Table 15 show that when the bias is below 3%, the performance degradation remains relatively minor compared to the bias level. However, once the bias reaches approximately 3% or higher, the error begins to increase more noticeably. This indicates that PDE-PFN maintains robustness against moderate systematic bias, but large bias eventually leads to a degradation in accuracy as expected.

## H.9 COMPUTATIONAL COST ANALYSIS RELATED TO THE NUMBER OF CONTEXT AND QUERY POINTS

In this section, we analyze how the number of context and query points affects the training and inference time as well as memory cost. For both training and inference, we report the average per-iteration runtime measured over 10 iterations and record the peak GPU memory usage. When evaluating the effect of context size, we fix the number of query points to 4,096. Likewise, when evaluating the effect of query size, we fix the number of context points to 4,096.

Table 16: Training and inference performance with different context lengths.

Context length	1,024	2,048	4,096	8,192	12,288	16,384
Training time (sec)	0.104	0.102	0.101	0.107	0.136	0.165
Training memory (MB)	280.0	329.8	420.7	617.2	800.4	987.4
Inference time (sec)	0.030	0.041	0.065	0.115	0.165	0.213
Inference memory (MB)	237.8	312.0	459.7	755.7	1054.7	1349.7

As shown in Tables 16 and 17, all four metrics increase linearly as the number of context or query points increases. This confirms that, due to our use of AGF layers, the model exhibits linear rather than quadratic complexity, unlike a *vanilla* Transformer.

1404  
1405  
1406  
1407  
1408  
1409  
1410  
1411  
1412  
1413  
1414  
1415  
1416  
1417  
1418  
1419  
1420  
1421  
1422  
1423  
1424  
1425  
1426  
1427  
1428  
1429  
1430  
1431  
1432  
1433  
1434  
1435  
1436  
1437  
1438  
1439  
1440  
1441  
1442  
1443  
1444  
1445  
1446  
1447  
1448  
1449  
1450  
1451  
1452  
1453  
1454  
1455  
1456  
1457  
Table 17: Training and inference performance with different **query lengths**.

Query length	1,024	2,048	4,096	8,192	12,288	16,384
Training time (sec)	0.101	0.100	0.101	0.123	0.148	0.176
Training memory (MB)	872.7	1010.1	1265.4	1788.6	2303.0	2831.2
Inference time (sec)	0.056	0.059	0.065	0.078	0.090	0.104
Inference memory (MB)	358.3	391.7	458.7	594.7	733.7	867.7

## H.10 2D CDR EQUATIONS WITH DIFFERENT INITIAL CONDITION

In this experiment, we examine whether our model continues to perform well on the 2D CDR equations from the main paper when the dataset is generated using a different initial condition. Specifically, we replace the original initial condition,  $1 + \sin(2x) \sin(2y)$ , with  $1 + \sin(3x) \sin(3y)$  to construct a new 2D CDR equation dataset. We then conduct training and evaluation following the same procedure used in the main experiments.

Table 18: The evaluation results for the spatiotemporal interpolation and temporal extrapolation tasks applied to the new family of 2D CDR equations. They are measured in the same way of Section 4.2.1. The best performance is marked in **bold** and the second-best performance is marked with underline.

Task	Metric	Ours	DeepONet	FNO	F-FNO	A-FNO	LNO	Poseidon	DPOT
Spatiotemporal interpolation	$L_1$ Abs	<b>0.01339</b>	<u>0.07394</u>	-	-	-	-	-	-
	$L_2$ Rel	<b>0.01885</b>	<u>0.09376</u>	-	-	-	-	-	-
	$L_\infty$ Rel	<b>0.06606</b>	<u>0.26320</u>	-	-	-	-	-	-
Temporal extrapolation	$L_1$ Abs	<b>0.02017</b>	0.03988	<u>0.03184</u>	0.03371	0.08342	0.04696	0.09397	0.12628
	$L_2$ Rel	<b>0.02929</b>	0.05282	<u>0.04160</u>	0.04192	0.09967	0.06282	0.10352	0.13633
	$L_\infty$ Rel	<b>0.10793</b>	0.17487	0.12671	<u>0.10980</u>	0.27716	0.18313	0.23855	0.13633

The experimental results in Table 18 show that our model continues to achieve the best performance even when evaluated on datasets generated with a new initial condition. This confirms that our model does not merely perform well on a single fixed initial condition, but instead generalizes robustly across different initial conditions. Furthermore, the experiments on the SWE and CNSE datasets in the main paper also support this conclusion, as both datasets consist of samples generated from varying initial conditions under the same underlying PDE. Thus, our model is not limited to datasets constructed with a single, unified initial condition.

## I EXPERIMENTAL DETAILS

### I.1 ENVIRONMENTS

The experiments on the family of 2D CDR equations, SWE, CNSE, and Darcy Flow were conducted using an NVIDIA RTX A6000, while the CNSE experiments were conducted on an NVIDIA RTX A5000. Details of the Python version and the packages used in the experiments can be found in the accompanying code’s environment file.

### I.2 COST COMPARISON ON THE MAIN EXPERIMENTS

In addition to the comparison between the baselines in Table 7, the additional comparison between the baselines is shown below in Table 19. In this tables, we compare the number of parameters, inference time, and the GPU memory usage for all models in their best settings. The inference time and GPU memory usage are measured in the test process on the family of 2D CDR equations temporal extrapolation task. In the experiment, both versions of our model are run with the same hyperparameters; therefore, we report the results under the single label Ours.

### I.3 HYPERPARAMETERS

In this section, we describe the number of data points for each equation and the hyperparameters used in the experiments for each model. In Table 20, the data points used in each experiment are shown separately for training, validation, and testing. For the training data, the context and query sets are constructed differently depending on the task. In the family of 2D CDR equations experiments

1458  
1459 Table 19: Additional comparisons between baselines and our model in the family of 2D CDR  
1460 equations temporal extrapolation test process.

Model	Number of parameters	inference time per sample(s)	GPU memory usage(MB)
Ours	4,488,873	0.24	26.64
DeepONet	723,969	0.01	11.06
FNO	143,268,097	0.10	1098.86
F-FNO	23,631,105	0.02	98.47
A-FNO	7,155,472	0.08	35.62
Poseidon	157,625,930	0.27	2408.40
DPOT	475,926,558	0.20	1825.72

1468  
1469 with unseen coefficients, we only performed testing without training or validation; therefore, only the  
1470 testing data points are reported.1471  
1472 Table 20: Number of data points used for training, validation, and testing in all experiments.

PDE	Task	Train	Valid		Test	
			Context	Query	Context	Query
Family of 2D CDR equations	Spatiotemporal interpolation	6,144	1,024	1,024	1,024	4,096
	Temporal extrapolation	6,144	1,024	1,024	1,024	4,096
	Unseen coeff temporal extrapolation	-	-	-	1,024	4,096
SWE	Spatiotemporal interpolation	98,304	16,384	16,384	16,384	65,536
	Temporal extrapolation	98,304	16,384	16,384	16,384	65,536
CNSE	Time trajectory predicting operator learning	20,480	4,096	16,384	4,096	16,384
Airfoil	Time trajectory predicting operator learning	26,165	5,233	20,932	5,233	20,932
Darcy Flow	$a \mapsto u$ operator learning	$a : 16,384$ $u : 16,384$				

1487  
1488 In the family of 2D CDR equation experiments, the evaluations on unseen coefficients are conducted  
1489 without additional training. Therefore, we report only the hyperparameters used for the spatiotemporal  
1490 interpolation and temporal extrapolation tasks trained on seen coefficients. During fine-tuning after  
1491 pre-training, the first linear layer of the encoder and the final linear layer of the decoder are re-  
1492 initialized to deal with cases where the number of variables to be predicted differs. For the family of  
1493 2D CDR equations, SWE, CNSE, and Airfoil experiments, we employ sinusoidal FFE, whereas FFT-  
1494 based embeddings are used in the Darcy Flow experiments. For fairness, all best hyperparameters were  
1495 selected based on those that achieved the highest performance on the validation. The hyperparameters  
1496 for our model and the baselines in each experiment can be found in Table from 21 to 23. In each  
1497 experiment, both versions of our model are run with the same hyperparameters; therefore, we report  
1498 the results under the single label Ours.

1512 Table 21: Best hyperparameter for each model used in the family of 2D CDR equations and SWE  
1513 experiments.

Model	Hyperparameter Name	Family of 2D CDR equations		SWE	
		Spatiotemporal interpolation	Temporal extrapolation	Spatiotemporal interpolation	Temporal extrapolation
Ours	Attention layers number			16	
	Attention hidden dimension			64	
	Attention head number			5	
	FFE dimension			0	
	AGF depth			15	
	En/Decoder hidden dimension			512	
DeepONet	Branch net depth	6	5	5	6
	Trunk net depth	5	5	6	6
	Hidden dimension	256	256	256	256
FNO	Layers number	-	4	-	3
	Hidden dimension	-	256	-	128
F-FNO	Layers number	-	5	-	3
	Hidden dimension	-	256	-	128
A-FNO	Layers number	-	12	-	3
	Hidden dimension	-	256	-	256
Poseidon	Pre-trained weight	-	L	-	T
DPOT	Pre-trained weight	-	L	-	M
PINN prior	Training loss threshold	$1 \times 10^{-4}$		-	
	Maximum training epoch	200		-	

1535 Table 22: Best hyperparameter for each model used in the CNSE and Darcy Flow experiments.

Model	Hyperparameter Name	CNSE	Darcy Flow
Ours	Attention layers number	16	
	Attention hidden dimension	64	
	Attention head number	5	
	FFE dimension	0	3
	AGF depth		15
	En/Decoder hidden dimension	512	
DeepONet	Branch net depth	6	6
	Trunk net depth	5	5
	Hidden dimension	256	256
FNO	Layers number	3	6
	Hidden dimension	256	128
F-FNO	Layers number	4	3
	Hidden dimension	256	256
A-FNO	Layers number	8	16
	Hidden dimension	256	128
Poseidon	pre-trained weight	T	T
DPOT	pre-trained weight	L	S

1554 Table 23: Best hyperparameter for each model used in the Airfoil experiments.

Model	Hyperparameter Name	Operator Learining
Ours	Attention layers number	16
	Attention hidden dimension	64
	Attention head number	5
	FFE dimension	2
	AGF depth	15
	En/Decoder hidden dimension	512
DeepONet	Branch net depth	6
	Trunk net depth	5
	Hidden dimension	256
Oformer	Layers number	3
	Hidden dimension	64