LEARNING DATA-EFFICIENT AND GENERALIZABLE NEURAL OPERATORS VIA FUNDAMENTAL PHYSICS KNOWLEDGE

Anonymous authorsPaper under double-blind review

ABSTRACT

Recent advances in scientific machine learning (SciML) have enabled neural operators (NOs) to serve as powerful surrogates for modeling the dynamic evolution of physical systems governed by partial differential equations (PDEs). While existing approaches focus primarily on learning simulations from the target PDE, they often overlook more **fundamental physical principles** underlying these equations. Inspired by how numerical solvers are compatible with simulations of different settings of PDEs, we propose a multiphysics training framework that jointly learns from both the original PDEs and their simplified basic forms. Our framework enhances data efficiency, reduces predictive errors, and improves outof-distribution (OOD) generalization, particularly in scenarios involving shifts of physical parameters and synthetic-to-real transfer. Our method is architectureagnostic and demonstrates consistent improvements in normalized root mean square error (nRMSE) across a wide range of 1D/2D/3D PDE problems. Through extensive experiments, we show that explicit incorporation of fundamental physics knowledge significantly strengthens the generalization ability of neural operators. We promise to release models and data upon acceptance.

1 Introduction

Recent advances in scientific machine learning (SciML) have broadened traditional machine learning (ML) for modeling physical systems, using deep neural networks (DNNs) especially neural operators (NOs) (Li et al., 2021a; Pathak et al., 2022; Lam et al., 2023; Bi et al., 2023) as fast, accurate surrogates for solving partial differential equations (PDEs) (Raissi et al., 2019; Edwards, 2022; Kochkov et al., 2021; Pathak et al., 2022). However, compared to numerical methods, **a key disadvantage** of recent data-driven SciML models is their **limited integration of fundamental physical knowledge**.



Figure 1: Can SciML models (e.g. neural operators trained on advanced PDEs) also understand fundamental physics knowledge (basic terms like diffusion, advection)?

Numerical solvers, though tailored to specific PDEs or discretizations, inherently preserve physical laws (e.g., conservation, symmetry), ensuring consistent and plausible simulations across diverse conditions (physical parameters, boundaries, geometries, etc.) (Ketcheson et al., 2012; Hansen et al., 2023; Mouli et al., 2024; Holl & Thuerey, 2024). In contrast, data-driven models, despite learning across PDE types (e.g., via multiphysics pretraining in SciML foundation models (McCabe et al., 2023; Hao et al., 2024)), remain sensitive to training distributions, degrading under distribution shifts (Subramanian et al., 2023; Benitez et al., 2024) and requiring large, diverse datasets. This fragility is worsened by the absence of rigorous verification: unlike classical solvers, SciML models are rarely evaluated against decomposed PDE components. This gap introduces three major challenges: 1) High data demands: Without physics priors, neural operators require large, diverse datasets for high precision, as seen in recent SciML foundation models(Hao et al., 2024; McCabe et al., 2023) which focus on generalization without addressing training data efficiency. 2) Physical inconsistency: Lacking inductive biases, these models may violate conservation laws or produce

unphysical outputs, particularly in long-term rollout predictions. 3) **Poor generalization**: Neural PDE solvers often struggle with unseen simulation settings and requires retraining.

Motivated by the above challenges, we ask two scientific questions:

Q1: Can neural operators understand both original PDEs and fundamental physics knowledge? Q2: Can neural operators benefit from explicit learning of fundamental physics knowledge?

In this paper, we highlight the importance of enforcing the learning of fundamental physical knowledge in neural operators. The key idea is to *identify physically plausible basic terms* that can be *decomposed from original PDEs*, and *incorporate their simulations during training*. Although often overlooked in SciML, our experiments demonstrate that these fundamental physical terms encode rich physical knowledge. Not only can they be utilized without incurring additional computational costs, but they also *widely offer substantial and multifaceted benefits*. This opens up a new door to improve the comprehensive generalization of neural operators with data efficiency.

We summarize our contributions below:

- 1. Through comprehensive evaluations of public SciML models, we observe a strong correlation between performance on original PDEs and basic PDE terms, highlighting the importance of fundamental physical knowledge in neural operators (Section 2.2).
- 2. We propose to explicitly incorporate fundamental physical knowledge into neural operators. Specifically, we design a simple and principled multiphysics strategy to train neural operators on simulations of both the original PDE and its basic form. (Section 3).
- 3. Our method exhibits three major benefits: 1) data efficiency (Section 4.2), 2) long-term physical consistency (Section 4.3), 3) strong generalization in OOD (Section 4.4) and real-world (Section 4.5) scenarios. We evaluate our method on a wide range of 1D/2D/3D PDEs, achieving consistent improvement in nRMSE (Section 4.2).

2 BACKGROUND

2.1 DEFINITION OF PDE LEARNING IN SCIML

For time-dependent PDEs, the solution is a vector-valued mapping $\boldsymbol{v}: \mathcal{T} \times \mathcal{S} \times \Theta \to \mathbb{R}^d$, defined on temporal domain \mathcal{T} , spatial domain \mathcal{S} , and parameter space Θ , with d the number of dependent variables. Numerical solvers compute $\boldsymbol{v}_{\theta}(t,\cdot)$ from $\ell \geq 1$ past steps, enabling finite-difference approximations: $\mathcal{N}_{\theta} = [\boldsymbol{v}_{\theta}(t-\ell \cdot \Delta t,\cdot),\ldots,\boldsymbol{v}_{\theta}(t-\Delta t,\cdot)] \mapsto \boldsymbol{v}_{\theta}(t,\cdot)$, where Δt is the temporal resolution. SciML aims to learn a surrogate operator $\widehat{\mathcal{N}}_{\theta,\phi}$, parameterized by physical parameters θ and learnable weights ϕ , that approximates this mapping. Given N simulations, $\mathcal{D} := \{\boldsymbol{v}^{(i)}([0:t_{\max}],\cdot) \mid i=1,\ldots,N\}$, the model is trained by minimizing a loss L, often the normalized root mean squared error: $nRMSE = \frac{\|\boldsymbol{v}_{pred} - \boldsymbol{v}\|_2}{\|\boldsymbol{v}\|_2}$, where \boldsymbol{v}_{pred} is the model prediction.

2.2 MOTIVATION: NEURAL OPERATORS EXHIBIT CORRELATED YET WORSE PERFORMANCE ON FUNDAMENTAL PHYSICS

We begin with a motivating example to highlight the importance of incorporating fundamental physical knowledge into neural operators. Specifically, we gather publicly released pretrained neural operators, with a focus on *SciML foundation models* that are jointly pretrained across multiple PDEs ("multiphysics") (McCabe et al., 2023; Hao et al., 2024). We evaluate these models on their ability to capture fundamental physical knowledge (formally defined in Section 3.1), which none of them were explicitly trained on, and compare their performance against the original PDE simulations.

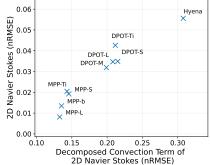


Figure 2: Neural operators and SciML foundation models (MPP (McCabe et al., 2023), DPOT (Hao et al., 2024), Hyena (Patil et al., 2023)) exhibit correlated yet worse performance on fundamental physics (2D incompressible Navier Stokes).

From Figure 2, we observe a strong Spearman correlation (0.967) between errors on original PDEs and their basic terms. This suggests that stronger SciML models implicitly learn basic PDE components more effectively. However, because these terms are not explicitly included in their training data, their absolute errors (0.133–0.308 on the x-axis) remain much larger than those for the original PDEs (0.008–0.056 on the y-axis). This gap reveals a key limitation: while the models demonstrate transfer across multiple physics, they still **fail to reliably capture the fundamental PDE components that underpin complex equations**. This motivates us to **explicitly** enforce an understanding of fundamental physical knowledge within neural operators.

3 METHODS

Motivated by our observation in Section 2.2, we incorporate fundamental physics knowledge into learning neural operators via: 1) Defining and decomposing basic forms from the original PDE (Section 3.1); 2) Jointly training neural operators on simulations from both basic forms and the original PDE (Section 3.2). We provide an overview of our method in Figure 3.

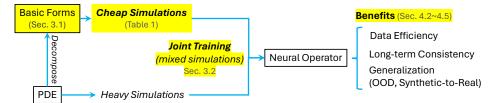


Figure 3: Overview of our method. Decomposed PDEs encode rich fundamental physical knowledge and introduce cheaper simulations. By jointly training on both the full PDE and its decomposed basic form, we bring multiple benefits to neural operators.

3.1 FUNDAMENTAL PHYSICAL KNOWLEDGE VIA DECOMPOSED BASIC PDE FORMS

3.1.1 DEFINING FUNDAMENTAL PHYSICAL KNOWLEDGE OF PDES

Let us consider the second-order PDE as a general example:

$$\sum_{i,j=1}^{n} a_{ij}(\boldsymbol{x},\boldsymbol{u},\nabla \boldsymbol{u}) \frac{\partial^{2} \boldsymbol{u}}{\partial x_{i} \partial x_{j}} + \sum_{i=1}^{n} b_{i}(\boldsymbol{x},\boldsymbol{u},\nabla \boldsymbol{u}) \frac{\partial \boldsymbol{u}}{\partial x_{i}} + c(\boldsymbol{x},\boldsymbol{u},\nabla \boldsymbol{u}) = f(\boldsymbol{x}), \tag{1}$$

where u is the target solution, with $x \in \mathbb{R}^n$ the physical space (e.g., n=3 for 2D time-dependent PDEs). The coefficients a_{ij}, b_i, c ("physical parameters") govern the dynamics; mismatches between training and testing values cause domain shifts, leading to out-of-distribution (OOD) simulations. Finally, f denotes an external forcing function (Nandakumaran & Datti, 2020).

We establish a systematic process to define the **fundamental physical knowledge of PDEs**, namely **basic PDE terms**: 1) retain terms that govern the essential and dominant physical dynamics; 2) remove terms that induce solver stiffness, increase computational cost, or contribute little to the pattern formation of interest. This procedure typically yields a simplified PDE form that can be simulated far more efficiently, while still capturing the key physical dynamics of the original system. From a machine learning perspective, decomposing a PDE into its basic form acts as a data augmentation strategy that reduces data collection costs.

Importantly, our approach of discovering and incorporating basic PDE forms differs fundamentally from the multiphysics training in recent SciML foundation models (McCabe et al., 2023; Hao et al., 2024), which simply aggregate diverse or even weakly related PDE systems. In contrast, we stress that neural operators must support basic PDE terms as a foundation while learning complex PDEs. Figure 4 illustrates the PDEs studied alongside their corresponding decomposed basic forms. An ablation study on different choices of terms is also provided in Appendix D.6.

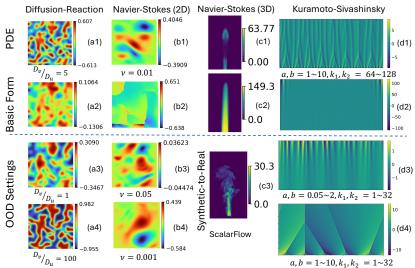


Figure 4: Visualizations of simulations of PDEs and their decomposed basic forms (Section 3.1). From left to right: Diffusion-Reaction (activator concentration), 2D Navier-Stokes (fluid velocity), 3D Navier-Stokes (smoke density), and Kuramoto-Sivashinsky (perturbation amplitude). Basic PDE forms are used for training neural operators with fundamental physics knowledge (Section 3.2), and the OOD settings are used for evaluating the generalization of neural operators (Section 4.4). D_v , D_u : diffusion coefficients (Equation 2). ν : viscosity (Equation 4). a, b, k_1 , k_2 : magnitudes and wavenumbers (Equation 7).

3.1.2 DIFFUSION-REACTION

The Diffusion-Reaction equation models an activator-inhibitor system, which typically happens in the dynamics of chemistry, biology, and ecology. The Diffusion-Reaction equation describes spatiotemporal dynamics where chemical species or biological agents diffuse through a medium and simultaneously undergo local reactions. These systems are often used to model pattern formation, such as Turing patterns, in domains ranging from chemistry to developmental biology.

$$\partial_t u = D_u \partial_{xx} u + D_u \partial_{yy} u + R_u, \quad \partial_t v = D_v \partial_{xx} v + D_v \partial_{yy} v + R_v. \tag{2}$$

In Equation 2, u and v represent the concentrations of activator and inhibitor, respectively, and D_u , D_v are diffusion coefficients. The nonlinear reaction terms R_u , R_v model biochemical interactions. In our experiments, we adopt the FitzHugh–Nagumo variant with $R_u(u,v) = u - u^3 - k - v$ and $R_v(u,v) = u - v$, where $k = 5 \times 10^{-3}$, consistent with values used in models of excitable media such as neurons or cardiac tissue.

Decomposed Basic Form. To isolate the fundamental transport behavior and reduce simulation cost, we consider a simplified form of Equation 2 by omitting the nonlinear reaction terms R_u and R_v , yielding pure diffusion equations:

$$\partial_t u = D_u \partial_{xx} u + D_u \partial_{yy} u, \quad \partial_t v = D_v \partial_{xx} v + D_v \partial_{yy} v. \tag{3}$$

- Why Drop Reaction Terms? This form retains the essential dispersal dynamics but eliminates the feedback coupling between u and v. Nonlinear reaction terms can vary rapidly, introducing stiffness into the PDE. This stiffness necessitates smaller time steps for stable numerical integration, increasing computational cost. By omitting these nonlinear terms, the system becomes linear and more amenable to efficient numerical solutions.
- Why Prioritize the Diffusion Term? Pure diffusion, though simpler, encodes key properties such as isotropic spreading and mass conservation, providing inductive bias for learning. Unlike reaction terms, which act locally to update activator and inhibitor concentrations, the diffusion term governs spatial coupling, as the primary source of pattern formation and spatial dynamics, facilitating transport and stabilization, which explains the visually similar patterns in Figure 4 a1 and a2.

Physical Scenarios. The emergence of spatial patterns in reaction-diffusion systems is governed by the ratio $\frac{D_v}{D_v}$, which affects the relative spreading rates of inhibitor and activator (Page et al.,

2005; Asgari et al., 2011; Gambino et al., 2024). Classical turing instability arises when $D_v \gg D_u$, leading to diffusion-driven pattern formation, and the inhibitor spreads out while the activator stays localized. Following previous works (Menou et al., 2023), we set $D_u = 1 \times 10^{-3}$, and focus on learning simulations when $\frac{D_v}{D_u} = 5$, and possible OOD scenarios when $\frac{D_v}{D_u} \in \{1, 100\}$.

3.1.3 INCOMPRESSIBLE NAVIER-STOKES

The Navier–Stokes equations govern the dynamics of fluid flow and serve as fundamentals for fluid simulations. It considers both the mass conservation and the momentum of fluid parcels.

$$\frac{\partial \boldsymbol{u}}{\partial t} = -(\boldsymbol{u} \cdot \nabla)\boldsymbol{u} + \nu \nabla^2 \boldsymbol{u} - \frac{1}{\rho} \nabla p + \boldsymbol{f}, \tag{4}$$

where u is the velocity field, ν is the dynamic viscosity, ρ is the fluid density, p is the fluid pressure, and f is the external force field.

Decomposed Basic Form. To isolate fundamental nonlinear transport mechanisms and reduce computational complexity, we simplify Equation 4 by omitting the pressure term (incompressibility via projection) $\frac{1}{\rho}\nabla p$ and diffusion term $\nu\nabla^2 u$:

$$\frac{\partial \boldsymbol{u}}{\partial t} = -(\boldsymbol{u} \cdot \nabla)\boldsymbol{u} + \boldsymbol{f},\tag{5}$$

This form captures inertial advection with external forcing and approximates high Reynolds number flows, where viscous effects are negligible. Such simplifications are analytically meaningful and are often used in turbulence modeling (e.g., inviscid Euler equations). Learning this reduced dynamics can help models internalize convection-dominant regimes.

- Why Drop Pressure and Diffusion Terms? The pressure term, which enforces fluid incompressibility, requires solving large linear systems and is difficult to parallelize and computationally expensive. Omitting it significantly accelerates the simulation. Similarly, the diffusion term in Navier-Stokes often uses explicit Euler integration with substeps, adding complexity. Removing it simplifies the simulation further. Moreover, for many visual effects like smoke or fire, viscosity is minimal, so the diffusion term has little visual impact and can often be omitted without noticeable loss in realism.
- Why Prioritize the Convection Term? Computationally, the convection term is cheap as it describes the local transport of fluid and no need to iterate across the spatial domain. Meanwhile, convection is the main driver of motion in most fluid flows, as it transports vorticity and mass. Without it, the fluid would just sit still or respond passively to forces. It captures nonlinear self-interaction, which is critical for dynamic, complex-looking behavior.

Physical Scenarios. In Navier-Stokes, the dynamic viscosity ν in Equation 4 (or the Reynolds number $Re = \frac{\rho u L}{\nu}$ where ρ is the density of the fluid, u is the flow speed, L is the characteristic linear dimension) controls the fluid dynamics. It measures the balance between inertial forces (pushes the fluid particles in different directions, leading to chaotic flow patterns) and viscous forces (resists motion and smoothes out differences in velocity, promoting an orderly flow) of a fluid. Following previous works (Schlichting & Gersten, 2016; Kochkov et al., 2021; Page et al., 2024), we will mainly focus on learning simulations when $\nu = 0.01$, and possible OOD scenarios when $\nu \in \{0.05, 0.001\}$, A smaller ν will lead to more turbulent flows.

3D Extension. In real-world scenarios such as atmospheric or smoke dynamics, buoyancy-driven flows provide additional complexity (Eckert et al., 2019). We extend our setting to simulate 3D incompressible Navier–Stokes in a rising plume scenario (see Figure 4 c1 and c3). We simulate how a plume of smoke rises and spreads in a 3D box-shaped environment. Smoke is introduced from a small circular inflow region located at the bottom center of the domain, at a steady inflow rate of 0.2 units per timestep. The smoke is carried upward due to buoyancy. This setting tests the method's robustness on complex spatiotemporal dynamics in three dimensions.

3.1.4 KURAMOTO-SIVASHINSKY

The Kuramoto-Sivashinsky equation is a nonlinear PDE that models the interplay of instability, nonlinearity, and dissipation, making it a prototype for studying spatiotemporal chaos. It is used to

simulate phenomena such as wrinkled flame fronts, thin fluid film instabilities, and chaotic pattern formation. We consider the following one-dimensional Kuramoto-Sivashinsky equation:

$$\frac{\partial \mathbf{u}}{\partial t} = -\mathbf{u}\partial_x \mathbf{u} - \partial_{xx} \mathbf{u} - \partial_{xxxx} \mathbf{u}, \quad \text{on } [0, L] \times [0, T]$$
(6)

The solution u generally represents the perturbation amplitude of a chaotic system. The spatial domain [0, L] is equipped with periodic boundary conditions. We impose the initial condition on $x \in [0, L]$ as

$$u_0(x) = a\cos\left(\frac{k_1\pi x}{L}\right) + b\cos\left(\frac{k_2\pi x}{L}\right) + \sigma\epsilon(x). \tag{7}$$

where u_0 is the superposition of two cosine modes and small mean-zero Gaussian perturbations. This initialization, common in KS studies (Papageorgiou & Smyrlis, 1991; Gudorf & Cvitanovic, 2019), combines deterministic cosine modes that inject controlled perturbations with small random noise. a,b,k_1,k_2 are tunable physical parameters controlling the strength and decay of the initial perturbation. Here, $\epsilon(x)$ are i.i.d. standard normal samples, and we set $\sigma=0.05$. Spatial derivatives are evaluated spectrally via discrete Fourier transforms, and time integration is performed with a fourth–order Runge–Kutta method.

Decomposed Basic Form. To isolate the linear stabilizing and destabilizing mechanisms and reduce computational complexity, we simplify the Kuramoto–Sivashinsky equation by omitting the nonlinear advection term $-u\partial_x u$. The reduced equation is

$$\frac{\partial \boldsymbol{u}}{\partial t} = -\partial_{xx}\boldsymbol{u} - \partial_{xxxx}\boldsymbol{u}, \quad \text{on } [0, L] \times [0, T], \tag{8}$$

which captures the competition between the destabilizing anti-diffusion term $(-\partial_{xx}u)$ and the stabilizing diffusion term $(-\partial_{xxx}u)$.

- Why Drop the Nonlinear Advection Term? For small amplitudes or short times, $-u\partial_x u$ is higher order in u, so the linearized dynamics govern the instabilities and pattern formation. Removing this term also eliminates costly Fourier transforms, yielding a significant speedup in simulation.
- Why Prioritize High-Order Stabilizing/Destabilizing Terms? The balance between destabilizing $-\partial_{xx}u$ and stabilizing $-\partial_{xxx}u$ determines how fast the chaotic model grows or decays. The fourth-order dissipation is particularly important for controlling stiffness and ensuring smoothness of solutions, making it critical for stable numerical integration. Finally, in practice, many interface/turbulence models reduce to precisely this "anti-diffusion + diffusion" structure, so analyzing these terms provides broadly transferable insights.

Physical Scenarios. In the initial condition (Equation 7), the amplitudes a, b control the strength of nonlinearity, while wavenumbers k_1, k_2 control the growth or decay of chaos, yielding regimes (Trefethen et al., 2001; Dawes, 2006): 1) For large a, b and high k_1, k_2 , the dynamics are predominantly **linear**, and chaos decays over time. 2) For small a, b and small k_1, k_2 , the system develops **weakly nonlinear** chaos, with instabilities growing slowly. 3) For large a, b and small k_1, k_2 , the system exhibits **strongly nonlinear** chaotic growth. In our experiments, we focus on learning in the linear regime with $a, b \in [1, 10]$ and $k_1, k_2 \in [64, 128]$, and evaluate on two out-of-distribution (OOD) regimes: weak nonlinearity $(a, b \in [0.05, 0.2], k_1, k_2 \in [1, 32])$ and strong nonlinearity $(a, b \in [1, 10], k_1, k_2 \in [1, 32])$.

3.2 JOINT LEARNING WITH FUNDAMENTAL PHYSICAL KNOWLEDGE

After defining our fundamental physics knowledge, we now explain how to integrate it into learning neural operators in principle from two perspectives data composition and neural architecture.

1) **Data Composition.** We jointly train neural operators on simulations of both our PDE and the decomposed basic form as a multiphysics training with a composite dataset. We summarize our simulations in Table 1. Since the simulation costs of decomposed basic forms are much cheaper than the original PDE, we can "trade-in" simulations of the original PDE for more simulations of basic forms **under the same simulation costs**. We define the "Sample Mixture Ratio" as the rate **derived**

Table 1: Summary of simulations of PDEs and their decomposed basic forms. GPU: NVIDIA RTX 6000 Ada.

PDE	Spatial Resolution	Temporal Steps	Target Variables	Simulation Costs (sec. per step)	Sample Mixture Ratio (PDE : Basic Form)
Diffusion-Reaction (Eq. 2) Basic Form (Eq. 3)	128×128	100	Activator $oldsymbol{u}$, Inhibitor $oldsymbol{v}$	$\begin{array}{c} 1.864 \times 10^{-2} \\ 6.610 \times 10^{-3} \end{array}$	1:3
Navier-Stokes (2D) (Eq. 4) Basic Form (2D) (Eq. 5)	256×256	1000	Velocity u , Density s	2.775 0.113	1:24
Navier-Stokes (3D) (Eq. 4) Basic Form (3D) (Eq. 5)	$50 \times 50 \times 89$	150	Velocity u , Density s	1.047 0.300	1:3
Kuramoto-Sivashinsky (Eq. 6) Basic Form (Eq. 8)	512	200	Perturbation Amplitude $oldsymbol{u}$	$\begin{array}{c} 1.176 \times 10^{-4} \\ 9.135 \times 10^{-6} \end{array}$	1:12

from the rate of simulation costs of original PDE with its decomposed basic form while making sure that "exchange" one primary data with the those number of basic form data will **maintain a comparable or reduced simulation budget when training baseline and our proposed model**.

We apply a multi-task formulation, where the model learns from both the original PDE and its simplified basic form. The idea is inspired from curriculum learning (Bengio et al., 2009; Pentina et al., 2014) and auxiliary task learning (Liu et al., 2019). In our method, the basic forms act as a simpler, physically motivated auxiliary task that can facilitate more efficient representation learning and accelerate the convergence on the primary task.

2) Neural Architecture. We mainly consider Fourier Neural Operator (FNO) (Li et al., 2021a) as our neural architecture. However, we make our method agnostic to specific architectures of neural operators: We generally share the backbone of the neural operator for learning both the main PDE and its basic terms, while employing separate final prediction layers for the two tasks. We will discuss the details on model structure in Appendix B, and will include more results using transformer (Dosovitskiy et al., 2020; Tong et al., 2022; McCabe et al., 2023; Chen et al., 2024) in Appendix D.

4 EXPERIMENTS

4.1 SETTINGS

Our baseline is learning the original PDE problem. In general, our method reallocates half of the baseline's simulation budget to simulate the basic PDE terms, with the sample mixture ratio defined in Table 1. To fairly compare with the baseline, we adopt the same **hyperparameters and optimization costs** (**number of gradient descent steps**). Since our goal is to evaluate performance on the original PDE, we use data from the basic PDE term only during training. All testing is conducted exclusively on the original PDE data with 100 samples (Takamoto et al., 2022). We use Adam optimizer, cosine annealing learning rate scheduler, and nRMSE defined in Section 2.1. We summarize our training details in Appendix C.

4.2 Data Efficiency

We first study our method with different numbers of training samples, and demonstrate that neural operators trained with our method can achieve stronger performance with less training data. We consider the following methods:

- "Baseline": Neural operators that are only trained on simulations of the original PDE.
- Ours: As described in Section 3.2, we can replace simulations of the original PDE with its decomposed basic form, allowing the total sim-

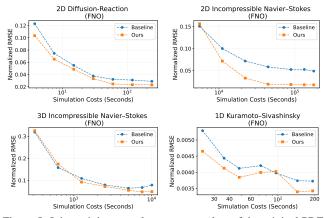


Figure 5: Joint training neural operators on data of the original PDE and the basic form improves performance and data efficiency. Y-axis: normalized RMSE. X-axis: simulation costs (seconds).

ulation cost of the training data can be comparable or even reduced. See Table 1 for the sample mixture rate for fair comparison.

In Figure 5, we study prediction errors of neural operators trained with different numbers of simulations (measured in their simulation costs). We can see that, across PDEs and neural architectures, our method (orange square) is to the lower left of the baseline (blue circle), which means that we can achieve improved prediction errors with reduced simulation costs.

385 386 387

389

390

391

392

394

395

396

397

398

399

400

401

402

403

404 405

406 407

408

409

410

411

412

413

426 427

428

429

430

431

LONG-TERM PHYSICAL CONSISTENCY 4.3

Next-frame prediction (Takamoto et al., 2022; McCabe et al., 2023; Hao et al., 2024) is a widely adopted evaluation, where input frames are always ground truth. Meanwhile, autoregressive inference, where model keeps rolling out to further temporal steps with its own output as inputs, is a meaningful and more challenging stress test. In autoregressive inference, a model forecasts futures with its own (noisy) output as inputs, and thus prediction error will accumulate along rollout steps. In our Figure 6, losses will be aggregated for five consecutive time steps. We can see that our improvements in Figure 5 further persist across autoregressive

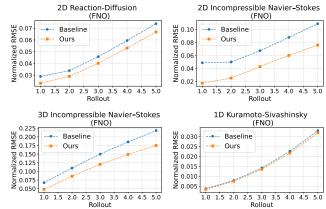


Figure 6: Joint training neural operators on data of the original PDE and the basic form improves performance with autoregressive inference at different unrolled steps. Models are evaluated using the best-performing checkpoints from training shown in Figure 5.

steps, leading to improved long-term consistency.

4.4 Out-of-Distribution (OOD) Generalization

We next show the benefits of our method towards the generalization of neural operators in out-ofdistribution (OOD) settings, where the physical parameters used during simulation are significantly shifted. We consider physical scenarios described in Sectin 3.1, and show results in Table 2. We can see that our method not only improves in-distribution errors, but also generalizes better on unseen physical dynamics (simulations by unseen parameters), leading to more robust neural operators.

Table 2: Comparisons of OOD generalization for different training methods. Models are evaluated using the best-performing checkpoints from training, as shown in Figure 5, under comparable simulation cost settings.

PDE	Model	Source	Target 1			Target 2	
		Setting	nRMSE	Setting	nRMSE	Setting	nRMSE
Diffusion-Reaction (2D)	Baseline	$\frac{D_v}{D_w} = 5$	0.0289	$\frac{D_v}{D_v} = 1$	0.0413	$\frac{D_v}{D_v} = 100$	0.0770
Diffusion Reaction (2D)	Ours	D_u	0.0231	D_u	0.0331	D_u	0.0538
Navier-Stokes (2D)	Baseline	$\nu = 0.01$	0.0487	$\nu = 0.05$	0.0825	$\nu = 0.0001$	0.0369
	Ours		0.0175		0.0222		0.0125
Navier-Stokes (3D)	Baseline	$\nu = 0.01$	0.0675	$\nu = 0.1$	0.0393	$\nu = 0.0001$	0.0836
Therefores (SD)	Ours	. , , , , , , , , , , , , , , , , , , ,	0.0481	2 0.1	0.0329	0.0001	0.0602
Kuramoto-Sivashinsky (1D)	Baseline	a, b = $1 \sim 10$	0.0037	a, b = $0.05 \sim 2$	0.0021	a, b = $1 \sim 10$	0.0200
	Ours	$k_1, k_2 = 64 \sim 128$	0.0034	$k_1, k_2 = 1 \sim 32$	0.0018	$k_1, k_2 = 1 \sim 32$	0.0197

4.5 SYNTHETIC-TO-REAL GENERALIZATION

Finally, we test neural operators trained on simulations of 3D Navier-Stokes in real-world scenarios. Essentially, transfering models trained on simulations to real observations is a synthetic-to-real generalization problem (Chen et al., 2020; 2021), as domain gaps between numerical simulations and real-world measurements always persist. We study the ScalarFlow dataset (Eckert et al., 2019), which is a reconstruction of real-world smoke plumes and assembles the first large-scale dataset of realistic turbulent flows. We provide visualizations of synthetic and ScalarFlow data in Figure 4.

We show the results and visualize the ground truth as well as the model predictions on smoke plumes from ScalarFlow in Figure 7. We can see that our method outperforms the baseline model and presents a qualitative comparison of scalar flow predictions on real data, illustrating that our jointly trained model exhibits improved synthetic-to-real generalization performance. Please read our Appendix D.2 for more results on 3D Navier Stokes.

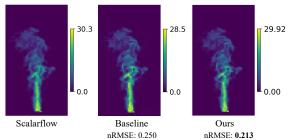


Figure 7: Visualizations of the last time step in the ScalarFlow and its predictions derived by baseline and our model.

5 RELATED WORKS

Machine Learning for Scientific Modeling Learning-based methods have long been used to model physical phenomena (Lagaris et al., 1998; 2000; Chen & Chen, 1995b;a). Physics-informed neural networks (PINNs) (Raissi et al., 2019; Zhu et al., 2019; Geneva & Zabaras, 2020; Gao et al., 2021; Ren et al., 2022) incorporate PDEs into loss functions to enforce physical laws, but often struggle with generalization and optimization issues (Krishnapriyan et al., 2021; Edwards, 2022). Operator learning methods like Fourier Neural Operators (FNO) (Li et al., 2021a; 2020; Kovachki et al., 2023) and DeepONet (Lu et al., 2019) offer greater flexibility by learning mappings between function spaces but require extensive labeled data (Raissi et al., 2019; Brandstetter et al., 2022; Zhang et al., 2023). We adopt a unique approach by evaluating and enhancing SciML through the lens of its compatibility with fundamental physical principles.

Data-Driven Neural PDE Solvers Machine learning has increasingly been used to approximate PDE solutions, with neural PDE solvers trained on diverse scenarios to mimic traditional simulators. Early models used CNNs (Guo et al., 2016; Zhu & Zabaras, 2018; Bhatnagar et al., 2019), while DeepONet (Lu et al., 2019) introduced a neural operator (NO) framework separating input and query encodings, inspiring many extensions (Wang et al., 2021; Hadorn, 2022; Wang et al., 2022; Lin et al., 2023; Venturi & Casey, 2023; Xu et al., 2023; McCabe et al., 2023; Hao et al., 2024). Advances like FNO (Li et al., 2021a), LNO (Cao et al., 2023), CNO (Raonic et al., 2024), and KNO (Xiong et al., 2024) have expanded the field, with FNO particularly impactful across applications (Li et al., 2021b; Guibas et al., 2021; Yang et al., 2021; Rahman et al., 2022b;a; Pathak et al., 2022; Liu et al., 2022). Compared with previous neural operator works, instead of naively swapping in cheaper simulations of simplified PDEs, the core merit of our work is to emphasize the multifaceted benefits of explicit learning of fundamental physics knowledge during operator learning.

Out-of-Distribution Generalization in SciML Interest in out-of-distribution (OOD) generalization for scientific machine learning (SciML) has grown recently. Subramanian et al. (2023) showed that fine-tuning neural operators (NOs) on OOD PDEs often requires many OOD simulations, which may be impractical. Benitez et al. (2024) proposed a Helmholtz-specific FNO with strong OOD performance, supported by Rademacher complexity and a novel risk bound. Other work includes ensemble methods leveraging uncertainty (Hansen et al., 2023; Mouli et al., 2024), loss functions informed by numerical schemes (Kim & Kang, 2024), and meta-learning for varied geometries (Liu et al., 2023). However, varying PDE types and setups across studies hinder unified insights into OOD generalization for NOs. Our work demonstrates that neural operators explicitly trained with fundamental physics knowledge exhibit improved OOD and synthetic-to-real generalization.

6 Conclusion

We present a principle and architecture-agnostic approach to enhance neural operators by explicitly incorporating fundamental physical knowledge into their training. By decomposing complex PDEs into simpler, physically meaningful basic forms and using them as auxiliary training signals, our proposed method significantly improves data efficiency, long-term predictive consistency, and outha ha h-of-distribution generalization. These improvements are demonstrated across a variety of PDE systems and neural operator architectures. Our finding highlights the untapped potential of fundamental physics as an inductive bias in scientific machine learning, offering a robust and cost-effective pathway to more reliable and generalizable surrogate models in real-world physical simulations.

THE USE OF LARGE LANGUAGE MODELS (LLMS)

LLMs did *not* play a significant role in either the research ideation or the writing of this paper. Their use was limited to correcting minor grammatical issues and typographical errors.

REFERENCES

- Yazdan Asgari, Mehrdad Ghaemi, and Mohammad Ghasem Mahjani. Pattern formation of the fitzhugh-nagumo model: cellular automata approach. 2011.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th Annual International Conference on Machine Learning*, ICML '09, pp. 41–48, New York, NY, USA, 2009. Association for Computing Machinery. ISBN 9781605585161. doi: 10.1145/1553374.1553380. URL https://doi.org/10.1145/1553374.1553380.
- Jose Antonio Lara Benitez, Takashi Furuya, Florian Faucher, Anastasis Kratsios, Xavier Tricoche, and Maarten V de Hoop. Out-of-distributional risk bounds for neural operators with applications to the helmholtz equation. *Journal of Computational Physics*, pp. 113168, 2024.
- Saakaar Bhatnagar, Yaser Afshar, Shaowu Pan, Karthik Duraisamy, and Shailendra Kaushik. Prediction of aerodynamic flow fields using convolutional neural networks. *Computational Mechanics*, 64:525–545, 2019.
- Kaifeng Bi, Lingxi Xie, Hengheng Zhang, Xin Chen, Xiaotao Gu, and Qi Tian. Accurate mediumrange global weather forecasting with 3d neural networks. *Nature*, 619(7970):533–538, 2023.
- Johannes Brandstetter, Max Welling, and Daniel E Worrall. Lie point symmetry data augmentation for neural pde solvers. In *International Conference on Machine Learning*, pp. 2241–2256. PMLR, 2022.
- Qianying Cao, Somdatta Goswami, and George Em Karniadakis. Lno: Laplace neural operator for solving differential equations. *arXiv* preprint arXiv:2303.10528, 2023.
- Tianping Chen and Hong Chen. Approximation capability to functions of several variables, nonlinear functionals, and operators by radial basis function neural networks. *IEEE Transactions on Neural Networks*, 6(4):904–910, 1995a.
- Tianping Chen and Hong Chen. Universal approximation to nonlinear operators by neural networks with arbitrary activation functions and its application to dynamical systems. *IEEE transactions on neural networks*, 6(4):911–917, 1995b.
- Wuyang Chen, Zhiding Yu, Zhangyang Wang, and Animashree Anandkumar. Automated synthetic-to-real generalization. In *International conference on machine learning*, pp. 1746–1756. PMLR, 2020.
- Wuyang Chen, Zhiding Yu, Shalini De Mello, Sifei Liu, Jose M Alvarez, Zhangyang Wang, and Anima Anandkumar. Contrastive syn-to-real generalization. *arXiv preprint arXiv:2104.02290*, 2021.
- Wuyang Chen, Jialin Song, Pu Ren, Shashank Subramanian, Dmitriy Morozov, and Michael W Mahoney. Data-efficient operator learning via unsupervised pretraining and in-context learning. *arXiv* preprint arXiv:2402.15734, 2024.
- Michael Crawshaw. Multi-task learning with deep neural networks: A survey, 2020. URL https://arxiv.org/abs/2009.09796.
- J. H. P. Dawes. Bifurcations and instabilities in dissipative systems: The kuramoto-sivashinsky equation (lecture notes). https://people.bath.ac.uk/jhpd20/lecturenotes/ks.pdf, 2006. Lecture note, March 21, 2006.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.

- Marie-Lena Eckert, Kiwon Um, and Nils Thuerey. Scalarflow: a large-scale volumetric data set of real-world scalar transport flows for computer animation and machine learning. *ACM Transactions on Graphics (TOG)*, 38(6):1–16, 2019.
 - C. Edwards. Neural networks learn to speed up simulations. *Communications of the ACM*, 65(5): 27–29, 2022.
 - G Gambino, MC Lombardo, R Rizzo, and M Sammartino. Excitable fitzhugh-nagumo model with cross-diffusion: close and far-from-equilibrium coherent structures. *Ricerche di Matematica*, 73 (Suppl 1):137–156, 2024.
 - Han Gao, Luning Sun, and Jian-Xun Wang. Phygeonet: Physics-informed geometry-adaptive convolutional neural networks for solving parameterized steady-state pdes on irregular domain. *Journal of Computational Physics*, 428:110079, 2021.
 - Nicholas Geneva and Nicholas Zabaras. Modeling the dynamics of pde systems with physics-constrained deep auto-regressive networks. *Journal of Computational Physics*, 403:109056, 2020.
 - Matthew Gudorf and Predrag Cvitanovic. Spatiotemporal tiling of the kuramoto-sivashinsky equation. In *APS March Meeting Abstracts*, volume 2019, pp. L70–263, 2019.
 - John Guibas, Morteza Mardani, Zongyi Li, Andrew Tao, Anima Anandkumar, and Bryan Catanzaro. Adaptive fourier neural operators: Efficient token mixers for transformers. *arXiv preprint arXiv:2111.13587*, 2021.
 - Xiaoxiao Guo, Wei Li, and Francesco Iorio. Convolutional neural networks for steady flow approximation. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 481–490, 2016.
 - Patrik Simon Hadorn. Shift-deeponet: Extending deep operator networks for discontinuous output functions. ETH Zurich, Seminar for Applied Mathematics, 2022.
 - Derek Hansen, Danielle C Maddix, Shima Alizadeh, Gaurav Gupta, and Michael W Mahoney. Learning physical models that can respect conservation laws. In *International Conference on Machine Learning*, pp. 12469–12510. PMLR, 2023.
 - Zhongkai Hao, Chang Su, Songming Liu, Julius Berner, Chengyang Ying, Hang Su, Anima Anandkumar, Jian Song, and Jun Zhu. Dpot: Auto-regressive denoising operator transformer for large-scale pde pre-training. *arXiv preprint arXiv:2403.03542*, 2024.
 - Philipp Holl and Nils Thuerey. Differentiable simulations for pytorch, tensorflow and jax. In *Forty-first International Conference on Machine Learning*, 2024.
 - David I Ketcheson, Kyle Mandli, Aron J Ahmadia, Amal Alghamdi, Manuel Quezada De Luna, Matteo Parsani, Matthew G Knepley, and Matthew Emmett. Pyclaw: Accessible, extensible, scalable tools for wave propagation problems. *SIAM Journal on Scientific Computing*, 34(4): C210–C231, 2012.
 - Taeyoung Kim and Myungjoo Kang. Approximating numerical fluxes using fourier neural operators for hyperbolic conservation laws. *CoRR*, 2024.
 - Dmitrii Kochkov, Jamie A Smith, Ayya Alieva, Qing Wang, Michael P Brenner, and Stephan Hoyer. Machine learning–accelerated computational fluid dynamics. *Proceedings of the National Academy of Sciences*, 118(21):e2101784118, 2021.
 - Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with applications to PDEs. *Journal of Machine Learning Research*, 24(89):1–97, 2023.
 - Aditi Krishnapriyan, Amir Gholami, Shandian Zhe, Robert Kirby, and Michael W Mahoney. Characterizing possible failure modes in physics-informed neural networks. *Advances in Neural Information Processing Systems*, 34:26548–26560, 2021.

Isaac E Lagaris, Aristidis Likas, and Dimitrios I Fotiadis. Artificial neural networks for solving ordinary and partial differential equations. *IEEE transactions on neural networks*, 9(5):987–1000, 1998.

- Isaac E Lagaris, Aristidis C Likas, and Dimitris G Papageorgiou. Neural-network methods for boundary value problems with irregular boundaries. *IEEE Transactions on Neural Networks*, 11 (5):1041–1049, 2000.
- Remi Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirnsberger, Meire Fortunato, Ferran Alet, Suman Ravuri, Timo Ewalds, Zach Eaton-Rosen, Weihua Hu, et al. Learning skillful medium-range global weather forecasting. *Science*, 382(6677):1416–1421, 2023.
- Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Andrew Stuart, Kaushik Bhattacharya, and Anima Anandkumar. Multipole graph neural operator for parametric partial differential equations. *Advances in Neural Information Processing Systems*, 33:6755–6766, 2020.
- Zongyi Li, Nikola Borislavov Kovachki, Kamyar Azizzadenesheli, Burigede liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021a.
- Zongyi Li, Hongkai Zheng, Nikola Kovachki, David Jin, Haoxuan Chen, Burigede Liu, Kamyar Azizzadenesheli, and Anima Anandkumar. Physics-informed neural operator for learning partial differential equations. *arXiv preprint arXiv:2111.03794*, 2021b.
- Guang Lin, Christian Moya, and Zecheng Zhang. B-deeponet: An enhanced bayesian deeponet for solving noisy parametric pdes using accelerated replica exchange sgld. *Journal of Computational Physics*, 473:111713, 2023.
- Burigede Liu, Nikola Kovachki, Zongyi Li, Kamyar Azizzadenesheli, Anima Anandkumar, Andrew M Stuart, and Kaushik Bhattacharya. A learning-based multiscale method and its application to inelastic impact problems. *Journal of the Mechanics and Physics of Solids*, 158:104668, 2022.
- Wenzhuo Liu, Mouadh Yagoubi, and Marc Schoenauer. Meta-learning for airflow simulations with graph neural networks. *arXiv preprint arXiv:2306.10624*, 2023.
- Yifan Liu, Bohan Zhuang, Chunhua Shen, Hao Chen, and Wei Yin. Auxiliary learning for deep multi-task learning, 2019. URL https://arxiv.org/abs/1909.02214.
- Lu Lu, Pengzhan Jin, and George Em Karniadakis. Deeponet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *arXiv* preprint arXiv:1910.03193, 2019.
- Michael McCabe, Bruno Régaldo-Saint Blancard, Liam Holden Parker, Ruben Ohana, Miles Cranmer, Alberto Bietti, Michael Eickenberg, Siavash Golkar, Geraud Krawezik, Francois Lanusse, et al. Multiple physics pretraining for physical surrogate models. *arXiv preprint arXiv:2310.02994*, 2023.
- Lucas Menou, Chengjie Luo, and David Zwicker. Physical interactions promote turing patterns. *arXiv preprint arXiv:2302.12521*, 2023.
- Grégoire Mialon, Quentin Garrido, Hannah Lawrence, Danyal Rehman, Yann LeCun, and Bobak T Kiani. Self-supervised learning with lie symmetries for partial differential equations. arXiv preprint arXiv:2307.05432, 2023.
- S Chandra Mouli, Danielle C Maddix, Shima Alizadeh, Gaurav Gupta, Andrew Stuart, Michael W Mahoney, and Yuyang Wang. Using uncertainty quantification to characterize and improve out-of-domain learning for pdes. *arXiv preprint arXiv:2403.10642*, 2024.
- AK Nandakumaran and PS Datti. *Partial differential equations: classical theory with a modern touch.* Cambridge University Press, 2020.
- Jacob Page, Peter Norgaard, Michael P Brenner, and Rich R Kerswell. Recurrent flow patterns as a basis for two-dimensional turbulence: Predicting statistics from structures. *Proceedings of the National Academy of Sciences*, 121(23):e2320007121, 2024.

- Karen M Page, Philip K Maini, and Nicholas AM Monk. Complex pattern formation in reaction—diffusion systems with spatially varying parameters. *Physica D: Nonlinear Phenomena*, 202(1-2): 95–115, 2005.
 - Demetrios T Papageorgiou and Yiorgos S Smyrlis. The route to chaos for the kuramoto-sivashinsky equation. *Theoretical and Computational Fluid Dynamics*, 3(1):15–42, 1991.
 - Jaideep Pathak, Shashank Subramanian, Peter Harrington, Sanjeev Raja, Ashesh Chattopadhyay, Morteza Mardani, Thorsten Kurth, David Hall, Zongyi Li, Kamyar Azizzadenesheli, et al. Fourcastnet: A global data-driven high-resolution weather model using adaptive fourier neural operators. arXiv preprint arXiv:2202.11214, 2022.
 - Saurabh Patil, Zijie Li, and Amir Barati Farimani. Hyena neural operator for partial differential equations, 2023. URL https://arxiv.org/abs/2306.16524.
 - Anastasia Pentina, Viktoriia Sharmanska, and Christoph H. Lampert. Curriculum learning of multiple tasks, 2014. URL https://arxiv.org/abs/1412.1353.
 - Md Ashiqur Rahman, Manuel A Florez, Anima Anandkumar, Zachary E Ross, and Kamyar Azizzadenesheli. Generative adversarial neural operators. *arXiv preprint arXiv:2205.03017*, 2022a.
 - Md Ashiqur Rahman, Zachary E Ross, and Kamyar Azizzadenesheli. U-no: U-shaped neural operators. *arXiv preprint arXiv:2204.11127*, 2022b.
 - Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019.
 - Bogdan Raonic, Roberto Molinaro, Tim De Ryck, Tobias Rohner, Francesca Bartolucci, Rima Alaifari, Siddhartha Mishra, and Emmanuel de Bézenac. Convolutional neural operators for robust and accurate learning of pdes. *Advances in Neural Information Processing Systems*, 36, 2024.
 - Pu Ren, Chengping Rao, Yang Liu, Jian-Xun Wang, and Hao Sun. Phycrnet: Physics-informed convolutional-recurrent network for solving spatiotemporal pdes. *Computer Methods in Applied Mechanics and Engineering*, 389:114399, 2022.
 - Sebastian Ruder. An overview of multi-task learning in deep neural networks, 2017. URL https://arxiv.org/abs/1706.05098.
 - Hermann Schlichting and Klaus Gersten. Boundary-layer theory. springer, 2016.
 - Shashank Subramanian, Peter Harrington, Kurt Keutzer, Wahid Bhimji, Dmitriy Morozov, Michael W Mahoney, and Amir Gholami. Towards foundation models for scientific machine learning: Characterizing scaling and transfer behavior. *arXiv preprint arXiv:2306.00258*, 2023.
 - Makoto Takamoto, Timothy Praditia, Raphael Leiteritz, Daniel MacKinlay, Francesco Alesiani, Dirk Pflüger, and Mathias Niepert. Pdebench: An extensive benchmark for scientific machine learning. *Advances in Neural Information Processing Systems*, 35:1596–1611, 2022.
 - Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *Advances in neural information processing systems*, 35:10078–10093, 2022.
 - Lloyd N. Trefethen, Kristine Embree, and contributors. *The (Unfinished) PDE Coffee Table Book*, 2001. URL https://people.maths.ox.ac.uk/trefethen/pdectb.html. Unpublished manuscript, freely available online.
 - Simone Venturi and Tiernan Casey. Svd perspectives for augmenting deeponet flexibility and interpretability. *Computer Methods in Applied Mechanics and Engineering*, 403:115718, 2023.
 - Sifan Wang, Hanwen Wang, and Paris Perdikaris. Learning the solution operator of parametric partial differential equations with physics-informed deeponets. *Science advances*, 7(40):eabi8605, 2021.
 - Sifan Wang, Hanwen Wang, and Paris Perdikaris. Improved architectures and training algorithms for deep operator networks. *Journal of Scientific Computing*, 92(2):35, 2022.

- Wei Xiong, Xiaomeng Huang, Ziyang Zhang, Ruixuan Deng, Pei Sun, and Yang Tian. Koopman neural operator as a mesh-free solver of non-linear partial differential equations. *Journal of Computational Physics*, pp. 113194, 2024.
- Wuzhe Xu, Yulong Lu, and Li Wang. Transfer learning enhanced deeponet for long-time prediction of evolution equations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 10629–10636, 2023.
- Yan Yang, Angela F Gao, Jorge C Castellanos, Zachary E Ross, Kamyar Azizzadenesheli, and Robert W Clayton. Seismic wave propagation and inversion with neural operators. *The Seismic Record*, 1(3):126–134, 2021.
- Xuan Zhang, Limei Wang, Jacob Helwig, Youzhi Luo, Cong Fu, Yaochen Xie, Meng Liu, Yuchao Lin, Zhao Xu, Keqiang Yan, et al. Artificial intelligence for science in quantum, atomistic, and continuum systems. *arXiv preprint arXiv:2307.08423*, 2023.
- Yinhao Zhu and Nicholas Zabaras. Bayesian deep convolutional encoder–decoder networks for surrogate modeling and uncertainty quantification. *Journal of Computational Physics*, 366:415–447, 2018.
- Yinhao Zhu, Nicholas Zabaras, Phaedon-Stelios Koutsourelakis, and Paris Perdikaris. Physics-constrained deep learning for high-dimensional surrogate modeling and uncertainty quantification without labeled data. *Journal of Computational Physics*, 394:56–81, 2019.

A SIMULATION SETTINGS

In this section, we detail the simulation settings for the 2D Diffusion-Reaction, 2D and 3D Incompressible Navier-Stokes and 1D Kuramoto-Sivashinsky (see Table 1 for the summary). We will also explain how we prepare simulations for the "Spatiotemporal Downsampling" in Section D.1.

To ensure a fair comparison under equivalent simulation costs with the basic form of each PDE, we downsample the original PDE simulations both spatially and temporally. We also introduce the settings for the corresponding reduced spatiotemporal resolution simulations. Note that the simulation cost of these downsampled settings is matched to that of the basic form, which implies that their sample mixture ratios in joint training remain equivalent.

To further explore the benefits of our multiphysics joint training approach with this reduced spatiotemporal resolution simulation strategies (refer to *Ours@Spatiotemporal* in Figure 9), we additionally introduce the simulation settings for the spatiotemporally downsampled basic form of the 2D Diffusion-Reaction equation. As what we will discuess in Section D.1, our framework is orthogonal to standard downsampling techniques, and combining the two can lead to further reductions in simulation cost. This reduction allows for an increased proportion of basic form samples in the training mixture under a fixed computational budget (see Table 3 for details).

A.1 DIFFUSION-REACTION

Our simulation setting for Diffusion-Reaction follows (Takamoto et al., 2022). Our solver is Py-Claw (Ketcheson et al., 2012) that uses the finite volume method. We set the initial condition as standard normal random noise $u(t=0,x,y) \sim \mathcal{N}(0,1.0)$. We use the homogeneous Neumann boundary condition. We simulate in a spatial domain of $\Omega = [-1,1]^2$, with resolution 128×128 . We simulate 5 seconds and save into 100 temporal steps.

Reduced Spatiotemporal Resolution. When simulating the original Diffusion-Reaction equation at low spatiotemporal grids (yellow curves in Figure 9), we reduce the spatial resolution from 128×128 to 96×96 , and reduce the number of temporal steps from 100 to 50. We then upsample to $128 \times 128 \times 100$ (steps) via bilinear interpolation to match the resolution of simulations of the original PDE. The total simulation interval is maintained at 5 seconds, preserving the underlying physical dynamics.

Similarly, we can further simulate our decomposed basic form of Diffusion-Reaction at low spatiotemporal resolution (green curve in Figure 9). Table 3 shows the simulation cost of decomposed basic forms with reduced spatiotemporal resolution and the sample mixture ratio.

Table 3: Summary of 2D Diffusion-Reaction simulations and its decomposed basic forms with reduced spatiotemporal resolution. "Sample Mixture Rate": We replace simulations of the original PDE with its decomposed basic form with reduced spatiotemporal resolution and make sure the total simulation cost of the training data can be comparable. GPU: NVIDIA RTX 6000 Ada.

PDE	Spatial Resolution	Temporal Steps	Target Variables	Simulation Costs (sec. per step)	Sample Mixture Ratio (PDE : Basic Form)
Diffusion-Reaction (Eq. 2)	128×128	100		1.864×10^{-2}	
Basic Form (Eq. 3) with Reduced Spatiotemporal Resolution	96×96	50	Activator u , Inhibitor v	2.390×10^{-3}	1:8

A.2 2D INCOMPRESSIBLE NAVIER-STOKES

Our simulation setting for incompressible Navier-Stokes follows (Takamoto et al., 2022). Our solver is PhiFlow (Holl & Thuerey, 2024). We simulate in a spatial domain of $\Omega = [0,1]^2$, with resolution 256×256 . We simulate 5 seconds with a $dt = 5 \times 10^{-5}$, and periodically save into 1000 temporal steps. Our initial conditions u_0 and forcing term f are drawn from isotropic Gaussian random fields, where the low-frequency components of the spectral density is scaled with scale and high-frequency components are suppressed with power-law decay by smoothness. For u_0 , scale is 0.15 and smoothness is 3.0. For f, scale is 0.4 and smoothness is 1.0. Boundary conditions are Dirichlet.

Reduced Spatiotemporal Resolution. When simulating the original 2D incompressible Navier-Stokes equation at low spatiotemporal grids (yellow curves in Figure 9), we reduce the spatial resolution from 256×256 to 100×100 . We then spatially upsample to 256×256 via bilinear interpolation to match the resolution of simulations of the original PDE. To reduce the temporal resolution while maintaining the same total simulation time and number of recorded frames, we increase the time-step size and proportionally reduce the number of integration steps and output interval. Specifically, we change the time-step from $dt = 5 \times 10^{-5}$ to $dt = 5 \times 10^{-4}$, and reduce the total number of time steps from $n_{\text{steps}} = 100,000$ to $n_{\text{steps}} = 10,000$. To preserve the temporal spacing between output frames, we decrease the frame interval from 100 to 10. This ensures the same total simulation duration of 5 seconds and the same number of output frames (1,000). This modification reduces computational cost by roughly 10 times.

A.3 3D INCOMPRESSIBLE NAVIER-STOKES

Our solver is PhiFlow (Holl & Thuerey, 2024). We simulate in a spatial domain of $\Omega = [0,1]^3$, with resolution $50 \times 50 \times 89$. We simulate 150 steps with a $dt = 2 \times 10^{-4}$. We set the initial u_0 as zero and upward buoyancy forcing term $f_z = 5 \times 10^{-4}$. Unlike the 2D Navier-Stokes, we introduce randomness of the buoyancy forcing term on horizontal directions by uniformly drawing f_x from $[-1,1] \times 10^{-4}$. We set Dirichlet zero boundary conditions.

A.4 1D KURAMOTO-SIVASHINSKY

Our simulation setting for the one-dimensional Kuramoto–Sivashinsky equation follows (Papageorgiou & Smyrlis, 1991; Gudorf & Cvitanovic, 2019). Our solver is a pseudospectral exponential time-differencing fourth-order Runge–Kutta scheme, implemented with 64 contour points. We use periodic boundary conditions on the spatial domain $\Omega = [0, L]$ with $L = 64\pi$, discretized with equispaced s = 512 grid points. Spatial derivatives are evaluated spectrally via FFT (fast fourier transform), and nonlinear terms are treated in physical space through pseudospectral squaring and differentiation. We simulate up to the final time T = 30, using an internal time step dt = 0.1, and record $n_{\rm steps} = 200$ snapshots uniformly in time.

B MODEL STRUCTURE

We mainly consider Fourier Neural Operator and transformer in this study. The model structures are shown in Figure 8. The basic form and the original PDE share all layers but the last prediction layer, which is agnostic to specific architectures of neural operators. This task-specific output layer is a very well-known and widely adopted configuration in multi-task learning. It has been highlighted in survey papers that hard parameter sharing (one input, shared hidden layers, multiple outputs) is the standard setup for multi-task learning due to its efficiency and representational benefits (Ruder, 2017; Crawshaw, 2020).

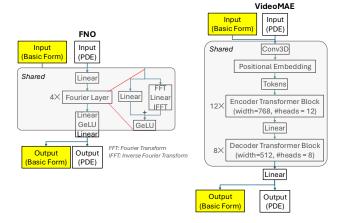


Figure 8: Our method is agnostic to specific architectures of neural operators: we always share the backbone of the model between learning of original PDE and its basic form, and decouple their predictions in the last layer.

C More Implementation Details

We summarize our training details in Table 4. We conducted our experiments on NVIDIA RTX 6000 Ada GPUs, each with 48 GB of memory.

Table 4: Training details. "DR": Diffusion-Reaction. "NS": Navier Stokes.

	2D DR (FNO)	2D DR (Transformer)	2D NS (FNO)	2D NS (Transformer)	3D NS (FNO)	3D NS (Transformer)	1D KS (FNO)
Input Shape Format	$H \times W \times$	$T \times C (C = 2)$	$H \times W \times T$	$\Gamma \times C (C = 3)$	$X \times Y \times Z$	\times T \times C (C = 4)	$T \times S$
Number of Training Samples (PDE Simulations)	2, 4, 8, 16	5, 32, 64, 128	2, 4, 8, 16	5, 32, 48, 64	2, 4, 8, 1	6, 32, 48, 64	$1024 \sim 8192$
Input Time Steps (\ell in Section 2.1)	10	10	10	10	10	10	10
Sample Mixture Ratio	1:3	1:3	1:24	1:6	1:3	1:3	1:12
Learning Rate	0.001	0.0003	0.001	0.001	0.001	0.00015	0.001
Batch Size for Primary Data	4	8	16	16	8	8	64
Epochs	100	60	20	30	20	80	50
Auxiliary Task Loss Weight	0.7	0.7	0.7	0.7	0.7	0.7	0.7
Training Hours	$0.08 \sim 1.83$	$0.6hr\sim7hr$	$1 \sim 29$	1.5~45	$0.5 \sim 6.5$	16~120	$2 \sim 18$
Gradient Descent Steps Per Epoch (Baseline and Ours)	$46\sim2912$	$23\sim1456$	$124\sim3960$	$124\sim3960$	$70\sim 2240$	$70\sim 2240$	$3040\sim24320$

D MORE RESULTS

Beyond the main comparison between the baseline and our proposed model using FNO showed in Section 4, we further conducted additional experiments to assess our approach. In this section, we will show the results from both FNO and Transformer.

D.1 COMPARISON OF DATA EFFICIENCY AND OUT-OF-DISTRIBUTION GENERALIZATION WITH SPATIOTEMPORAL DOWNSAMPLING

To demonstrate the effectiveness of our method, we conducted an ablation study comparing it with Spatiotemporal Downsampling on both 2D Diffusion–Reaction (DR) and 2D Navier–Stokes (NS). We defined the Spatiotemporal Downsampling method as follows: Neural operators that are trained with a mixture of simulations of the original PDE and simulations at low spatial and temporal resolutions (then linearly interpolated to the original resolution). Similar to our method, we can also save simulation costs with reduced spatiotemporal resolutions. See Appendix A for rates of downsampling and more details). Meanwhile, as our decomposed basic form is orthogonal to spatiotemporal downsampling during simulations, our method can serve as a complementary data augmentation. On 2D Diffusion-Reaction, we can simulate our decomposed basic forms at lower spatiotemporal resolution, leading to further reduced simulation costs and improved data efficiency (green triangle), outperforming the baseline at the lower spatiotemporal resolution (yellow diamond). For 2D Navier–Stokes with the Transformer, we adopt a 1:6 mixture ratio (vs. 1:24 for FNO) due to GPU memory limits; despite this disadvantage, our method still outperforms both the baseline and spatiotemporal downsampling at comparable cost.

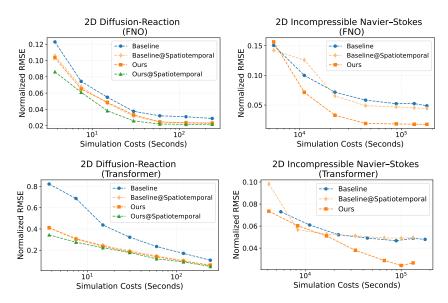


Figure 9: Joint training neural operators on data of the original PDE and the basic form improves performance and data efficiency. Rows: (top) FNO, (bottom) transformer. Columns: (left) 2D Diffusion-Reaction, (right) 2D Navier Stokes. "Spatiotemporal": short for "Spatiotemporal Downsampling". Y-axis: normalized RMSE. X-axis: simulation costs (seconds).

Table 5 reports the out-of-distribution (OOD) generalization results across both the 2D Diffusion-Reaction and Navier-Stokes equations. Similar to the results in Table 2, here we can see that our approach not only improves in-distribution errors but also consistently enhances generalization to simulations of unseen physical parameters. This robustness holds across both FNO and Transformer architectures, leading to more reliable and consistent neural operators under varying conditions.

Table 5: Comparisons of OOD generalization for different training methods with Transformer. Models are evaluated using the best checkpoints from training in Figure 5, under comparable simulation cost settings. "Spatiotemporal": short for "Spatiotemporal Downsampling".

PDE	Model	Sou	rce	Target 1		Target 2	
122	1,10001	Setting	nRMSE	Setting	nRMSE	Setting	nRMSE
	Baseline		0.0289	$\frac{D_v}{D_u} = 1$	0.0413	$\frac{D_v}{D_u} = 100$	0.0770
Diffusion-Reaction	Baseline@Spatiotemporal	$\frac{D_v}{D_u} = 5$	0.0234		0.0303		0.0663
(2D, FNO)	Ours	$D_u = 0$	0.0231		0.0331		0.0538
	Ours@Spatiotemporal		0.0218		0.0298		0.0596
	Baseline	$\frac{D_v}{D_u} = 5$	0.1056	$\frac{D_v}{D_u} = 1$	0.1249	$\frac{D_v}{D_u} = 100$	0.1976
Diffusion-Reaction (2D, Transformer)	Baseline@Spatiotemporal		0.0542		0.0698		0.0812
	Ours		0.0602		0.0782		0.0853
	Ours@Spatiotemporal		0.0469		0.0489		0.0671
	Baseline		0.0487		0.0825		0.0369
Navier-Stokes (2D, FNO)	Baseline@Spatiotemporal	$\nu = 0.01$	0.0442	$\nu = 0.05$	0.0743	$\nu = 0.0001$	0.0269
(2D, 1110)	Ours	•	0.0175		0.0222		0.0125
Navier-Stokes (2D, Transformer)	Baseline		0.0479		0.0853		0.0685
	Baseline@Spatiotemporal	$\nu = 0.01$	0.0496	$\nu = 0.05$	0.0568	$\nu = 0.0001$	0.0402
	Ours		0.0265		0.0397		0.0256

D.2 DATA EFFICIENCY AND OUT-OF-DISTRIBUTION GENERALIZATION OF TRANSFORMER FOR 3D NAVIER STOKES

Similar as what we have studied in Section 4, we aim to also demonstrate three key benefits: data efficient, long-term physical consistency, and strong generalization in OOD simulations using Transformer, for 3D Navier-Stokes as well.

In Figure 10, we can see that joint training (orange square) on both the original and basic forms of the 3D Navier-Stokes equation consistently reduces normalized RMSE from baseline (blue circle) across varying simulation budgets. This improvement is observed for Transformer architectures, highlighting enhanced data efficiency and generalization, which aligns with the results in Section 4.2.

In Table 6, we show that our joint training approach significantly improves out-of-distribution generalization on 3D Navier-Stokes across all test settings, outperforming the baseline for both FNO and Transformer models. Together with the results in Table 2 and 5, the consistent gains observed across all OOD setting results underscore the effectiveness and robustness of our method in generalizing to previously unseen physical regimes, particularly under significant shifts in simulation parameters.

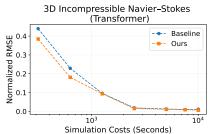


Figure 10: Joint training neural operators on data of the original 3D Navier-Stokes equation and the basic form improves performance and data efficiency.

Table 6: Comparisons of OOD generalization on 3D NS for different training methods using Transformer. Models are evaluated using the best checkpoints from training in Figure 10.

PDE	PDE Model		Source		Target 1		Target 2	
		Setting	nRMSE	Setting	nRMSE	Setting	nRMSE	
	Baseline		0.0114		0.0327		0.0816	
3D Navier-Stokes	Ours	$\nu = 0.01$	0.0064	$\nu = 0.1$	0.0124	$\nu = 0.0001$	0.0322	

D.3 More Long-Term Consistency Results of Transformer

In our Figure 11, we show the rollout performance of transformer on the 2D Diffusion-Reaction and 2D incompressible Navier-Stokes equation. Here, we run the experiments with the best checkpoints from training in Figure 9. Losses will be aggregated for five consecutive time steps. We can see that our improvements in Figure 9 further persist across autoregressive steps, leading to improved long-term consistency, aligning with our results in Figure 6.

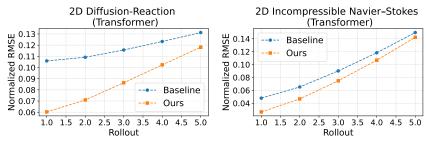


Figure 11: Joint training neural operators on data of the original PDE and the basic form improves performance with autoregressive inference at different unrolled steps using Transformer. Models are evaluated using the best-performing checkpoints from training shown in Figure 9.

D.4 MORE RANDOM SEEDS

To ensure the statistical robustness of our findings, we now run FNO using three different random seeds during initialization and training. For each configuration, we report the average performance across the three runs, and include standard deviation as error bars in all plots in Figure 12. This enables a more rigorous evaluation of model performance, capturing the inherent variance and mitigating the risk of overinterpretation from single-seed outcomes. We can see that the results demonstrate that joint training of neural operators on data from both the original PDE and its decomposed basic form yields consistent improvements in predictive performance and data efficiency, highlighting the effectiveness of this multiphysics learning strategy.

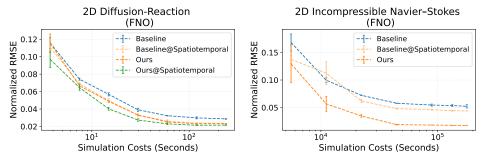


Figure 12: Model performance averaged over three random seeds. Joint training neural operators on data of the original PDE and the basic form consistently improves performance and data efficiency. Dash lines indicate the mean performance, with error bars representing standard deviation. Legends align with the descriptions in Section 4.2. Columns: (left) 2D Diffusion-Reaction, (right) 2D Navier Stokes. Y-axis: nRMSE. X-axis: Simulation Costs (seconds).

D.5 Loss Reweighting

In Section 3.2, we define our total loss for joint learning of the original PDE (Loss_{Full}) and its fundamental physical knowledge (decomposed basic form Loss_{Basic}) as

$$Loss = Loss_{Full} + 0.7 \times Loss_{Basic}$$

To address the concern regarding the fixed auxiliary loss weight, we conducted an ablation study on the effect of auxiliary loss weighting in joint training using FNO for the 2D Diffusion-Reaction system. We evaluate model performance across three auxiliary weight settings, 0.5, 0.7, and 1.0, by averaging accuracy over the range of training sample sizes (Figure 5) used in the data efficiency

experiments. In Table 7, we show the model averaged nRMSE, which is largely consistent with only minor variations in normalized RMSE. This result demonstrates that the model is largely insensitive to the specific choice of auxiliary weight and suggests that the improvements achieved through joint training are robust with respect to this hyperparameter. Based on this analysis, we fix the auxiliary weight to 0.7 for all the experiments.

Table 7: Ablation study on the auxiliary loss weight in joint training using FNO for the 2D Diffusion-Reaction across three settings Results of averaged nRMSE demonstrate consistent performance

Auxiliary Weights	0.5	0.7	1
Averaged nRMSE	0.0508	0.0491	0.0472

D.6 CHOICE OF FUNDAMENTAL TERM

 To demonstrate the importance of the choice of dropping terms, we conduct an ablation study on 2D Diffusion-Reaction, where we simulate reaction term instead of the fundamental diffusion term. The simulation cost for reaction-only term is 2.048×10^{-3} seconds per step, corresponding to a 1:9 sample mixture ratio when compared to the simulation cost of the original PDE data, making sure the comparable results.

From Figure 13, we can find that keeping the reaction term and removing the fundamental diffusion term will damage the accuracy with up to 64% increase of nRMSE, while applying the fundamental diffusion term keep boosting the model performance with 11% to 24% decrease of nRMSE. This ablation study confirms

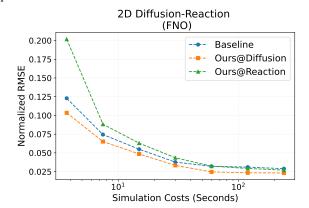


Figure 13: Ablation study of joint training neural operators on data of 2D Diffusion-Reaction equation with two different decomposed terms: the fundamental term - Diffusion, and Reaction shows the importance of choice on fundamental term from PDE equation. Y-axis: nRMSE. X-axis: Simulation Costs (Seconds).

that **the correct fundamental basic term** can improve the data-efficiency when joint training with the original data, and proved that the source of improvement in Figure 5 is clearly from training with the fundamental term itself.

D.7 LIE TRANSFORM ARGUMENT ON 2D IMCOMPRESSIBLE NAVIER STOKES

Lie symmetries offer a way to generate new, physically valid training examples by exploiting the analytic group transformations that map one PDE solution to another. This enables the model to learn representations that are inherently equivariant to fundamental symmetries such as translation, rotation, and scaling. To further prove the strength of our model, we leverage the implementation of Lie point symmetry augmentation from Brandstetter et al. (2022); Mialon et al. (2023), which is orthogonal to our multiphysics joint training approach, to 2D incompressible Navier Strokes equation.

We incorporate the augmentation process into our model. We only apply Lie-transform augmentations exclusively to the velocity (u) of the original 2D incompressible Navier-Stokes, leaving the remaining density and all target variables from the decomposed basic forms unchanged. Following (Mialon et al., 2023), the Lie transformation is implemented with a second-order Lie-Trotter splitting scheme with two steps, where the five fields (x,y,t,u_x,u_y) were transformed in accordance with the sampled generator strengths as follows: a maximum time shift (g_1) strength of 0.1, maximum spatial translations (g_2, g_3) strength of 0.1 in x and y respectively, a maximum scaling (g_4) strength of 0.05, a maximum rotation (g_5) strength of (g_5) streng

As our decomposed basic form is orthogonal to the Lie point symmetry augmentation, our method can serve as a complementary data augmentation. In Figure 14, we study prediction errors (nRMSE) of neural operators trained with different numbers of training samples (simulations). As we have already seen (Figure 5), our approach (orange square) significantly outperforms the baseline (blue circle). In contrast, the Lietransform augmentation alone (yellow diamond) only marginally improves the baseline. As a result, combining our approach with Lie transformations (green triangle) yields strong performance, but is compara-

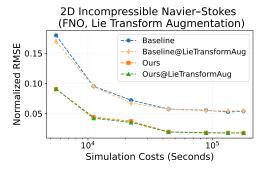


Figure 14: Joint training neural operators on data of the original PDE and the basic form, as a complementary data augmentation orthogonal to Lie-transform augumentation, can further improve performance and data efficiency. Y-axis: normalized RMSE. X-axis: simulation costs (seconds).

ble with our approach alone, underscoring the orthogonal and complementary benefits of these two techniques.

D.8 VISUALIZATION OF PREDICTIONS

To show the predicted PDE solution from our jointly training neural operators on original PDE equation and its basic form aligns with the ground truth, we present qualitative visualizations of model predictions across three PDEs, 2D Diffusion-Reaction, 2D and 3D Incompressible Navier-Stokes, and 1D Kuramoto–Sivashinsky, in Figure 15. For each PDE solution, we show the initial condition and predicted states at intermediate and final rollout times. The predictions are generated using the FNO model trained with our joint training framework. Across all systems and time points, the predictions closely align with the expected dynamics, accurately capturing both spatial patterns and temporal evolution. These visualizations highlight the model's capacity to generalize across scales and exhibit physically coherent behavior.

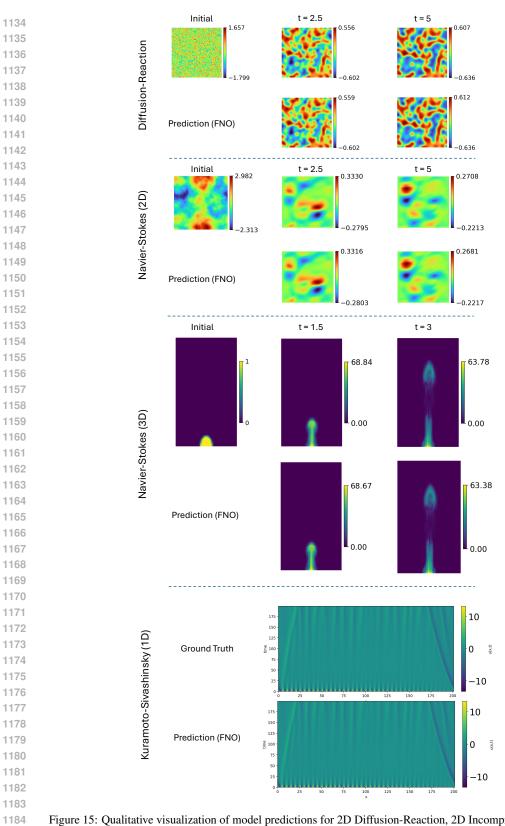


Figure 15: Qualitative visualization of model predictions for 2D Diffusion-Reaction, 2D Incompressible Navier-Stokes, 3D Incompressible Navier-Stokes systems and 1D Kuramoto–Sivashinsky using FNO trained with our joint framework. For the first three case, the initial state and predicted states at intermediate and final rollout times are shown. For 1D Kuramoto–Sivashinsky, we show the whole predicted trajectory here. The results demonstrate accurate temporal evolution and spatial coherence.