ACCELERATING DATA GENERATION FOR NONLINEAR TEMPORAL PDES VIA HOMOLOGOUS PERTURBATION IN SOLUTION SPACE

Anonymous authorsPaper under double-blind review

000

001

002

004

006

008 009 010

011 012

013

014

015

016

017

018

019

021

023

025

026

027

028

029

031

033 034

035

037

040

041

042

043

044

045

046 047

048

051

052

ABSTRACT

Data-driven deep learning methods like neural operators have advanced in solving nonlinear temporal partial differential equations (PDEs). However, these methods require large quantities of solution pairs—the solution functions and right-hand sides (RHS) of the equations. These pairs are typically generated via traditional numerical methods, which need thousands of time steps iterations far more than the dozens required for training, creating heavy computational and temporal overheads. To address these challenges, we propose a novel data generation algorithm, called HOmologous Perturbation in Solution Space (HOPSS), which directly generates training datasets with fewer time steps rather than following the traditional approach of generating large time steps datasets. This algorithm simultaneously accelerates dataset generation and preserves the approximate precision required for model training. Specifically, we first obtain a set of base solution functions from a reliable solver, usually with thousands of time steps, and then align them in time steps with training datasets by downsampling. Subsequently, we propose a "homologous perturbation" approach: by combining two solution functions (one as the primary function, the other as a homologous perturbation term scaled by a small scalar) with random noise, we efficiently generate comparable-precision PDE data points. Finally, using these data points, we compute the variation in the original equation's RHS to form new solution pairs. Theoretical and experimental results show HOPSS lowers time complexity. For example, on the Navier-Stokes equation, it generates 10,000 samples in approximately 10% of traditional methods' time, with comparable model training performance.

1 Introduction

Nonlinear temporal partial differential equations (temporal PDEs) serve as a core mathematical tool for precisely characterizing continuous physical systems in the real world that evolve dynamically over time. They possess exceptionally high application value across diverse real-world scenarios: for instance, the Navier-Stokes equations (Temam (2024)) describe the motion states of atmospheric fluids. Traditionally, solving PDEs has often relied on extensive domain expertise and computationally intensive numerical methods (e.g., the finite difference method Godunov & Bohachevsky (1959), finite element methods Strang et al. (1973)). With the rapid advancement of deep learning, a new physical dynamics modeling and prediction paradigm, such as neural operators, has sparked widespread discussion. Deep learning models can unearth latent physical relationships from data and predict future states at lower computational cost—a capability driving numerous breakthroughs in physical dynamics research, modeling, and prediction.

Neural operators—especially transformer-based ones—show promise for accelerating PDE solutions. But they rely heavily on large-scale training data, which comes from costly classical methods (e.g., FEMHughes (2000)). This creates a hard-to-solve circular dependency issue (Brandstetter et al. (2022)). This dependency becomes particularly problematic at industrial scales, as the high computational overhead of data generation via traditional numerical solvers severely hinders real-world applications (Zhang et al. (2025)). By accelerating dataset generation, we not only mitigate the high costs of classical method-driven model order reduction but also enhance sample complexity. Furthermore, efficient data generation supports broader generalization capabilities, which remains

an underexplored area critical for robust real-world applications(Kochkov et al. (2021), Stachenfeld et al. (2021)).

However, existing methods for generating PDE datasets exhibit notable limitations. Traditional numerical methods(e.g., the Crank-Nicolson method)typically require iterating over thousands of time steps to achieve solution convergence and stability, incurring significant temporal and computational overheads. While progress has been made in addressing linear time-independent PDEs (e.g., DiffOAS Dong et al. (2024), Wang et al. (2024)), these approaches rely on the linearity and time-invariance of the equations. This renders them incompatible with the data needs of neural solvers for nonlinear temporal PDE scenarios.

To address the challenges in forward modeling for nonlinear temporal PDEs, we propose a novel and efficient data generation algorithm termed HOmologous Perturbation in Solution Space (HOPSS). HOPSS simultaneously accelerates dataset generation while preserving the precision of the generated data. Specifically, HOPSS first generates a set of solution pairs under initial conditions, which serves as base solution functions in the solution space. These base solution functions are typically generated using a traditional solver to ensure high precision, and are then aligned to the time steps required for model training. Next, the base functions are fed into the generator, where two solutions are randomly selected: One is scaled by a small scalar, acting as the homologous perturbation term. It is then combined with the other to generate a new solution, plus small random noise. Concurrently, HOPSS computes the right-hand side (RHS) of the physical equation, ensuring compliance with the properties of the corresponding physical equation. A key advantage of HOPSS lies in its ability to avoid extensive iteration over a large number of time steps; instead, it acts directly on the time steps needed for training.

The distinct contributions of our work can be summarized as follows.

- We propose a novel data generation algorithm (HOPSS) tailored for nonlinear temporal PDEs. This algorithm simultaneously accelerates dataset generation while preserving the precision of the generated data—enabling the generation of large-scale nonlinear temporal PDE datasets at affordable time costs and computing resources.
- We demonstrate that our proposed algorithm significantly lowers computational complexity
 and shortens data generation time when solving nonlinear temporal PDEs. Notably, even
 with only approximately 10% of the generation time required by existing methods, neural operators trained on HOPSS-generated data exhibit performance comparable to those
 trained on data from conventional generation approaches.

2 RELATED WORK

2.1 Data-Driven Deep Learning for Solving temporal PDEs

Data-driven deep learning has become a transformative force in solving temporal PDEs, with key advancements centered on efficient model architectures and hybrid computational paradigms. Neural operators stand out as a major breakthrough: the Fourier Neural Operator (FNOLi et al. (2021)) and Deep Operator Network (DeepONetLu et al. (2021))leverage deep learning to capture complex spatiotemporal dependencies in PDE systems, outperforming conventional methods in efficiency for time-evolving problems. Complementary efforts target bottlenecks in classical PDE solving—for instance, studies exploring neural networks to accelerate linear equation system solutions directly reduce the computational overhead of temporal PDE workflows. Additionally, data-driven solvers have advanced via hybrid designs: works like Hsieh et al. (2019), Yang et al. (2016), and Kochkov et al. (2021) propose data-optimized iterative schemes, merging machine learning with traditional numerical techniques to boost efficiency for temporal PDE scenarios. Together, these directions highlight a shift toward data-centric methods that address the unique spatiotemporal challenges of temporal PDEs.

2.2 Data Generation for PDE Algorithms

Training data-driven PDE algorithms, particularly those targeting nonlinear temporal PDEs, demands large-scale offline paired parametrized datasets. These datasets must capture complex spatiotemporal dynamics—including nonlinear interactions (e.g., convection terms in Navier-Stokes

equations) and transient behaviors—making their generation computationally intensive. Traditionally, such datasets are produced exclusively via classical computational mathematics algorithms, following a standardized numerical workflow tailored to solve nonlinear temporal PDEs.

Numerically solving nonlinear temporal PDEs relies on two-step discretization, which converts continuous nonlinear equations into a sequence of solvable linear (or linearized) systems (Morton & Mayers (2005)). First, temporal discretization splits the PDE's long-term evolution into discrete-step iterative updates: explicit schemes (e.g., 4th-order Runge-Kutta, RK4 Hairer et al. (1993)) work for stable, low-stiffness systems, while implicit schemes (e.g., Crank-Nicolson Crank & Nicolson) are favored for stiff nonlinear systems to bypass strict time step constraints. Second, spatial discretization approximates spatial derivatives via core techniques (e.g., Finite Difference, Finite Element, Finite Volume Methods Strikwerda (2004); Hughes (2000); LeVeque (2002)), each tailored to specific needs like regular domains or conservation-focused transport PDEs. Notably, nonlinear terms (e.g., quadratic velocity in fluid dynamics) require linearization (e.g., Newton-Raphson iteration) to produce linear systems at each time step.

While traditional numerical methods for generating nonlinear temporal PDE datasets are highly mature, they demand substantial time and computational resources—largely due to the repeated solution of multiple matrix multiplications at each time step (Hao et al. (2022)). This inefficiency poses a significant bottleneck to the advancement of data-driven deep learning algorithms for PDEs, as large-scale dataset generation often becomes prohibitively costly. To address this challenge, emerging data generation methods have been proposed, with DiffOAS (Dong et al. (2024)), Wang et al. (2024), and Brandstetter et al. (2022) being representative examples. However, DiffOAS is specifically designed for data generation of linear time-independent PDEs; it cannot be directly extended to nonlinear temporal PDEs. A critical limitation is that DiffOAS relies heavily on the linear properties and steady-state characteristics of the target equations. This dependence renders it inadequate for nonlinear temporal scenarios, as it fails to preserve the spatiotemporal dynamics and physical precision required for training reliable data-driven solvers (Dong et al. (2024)).

3 Preliminaries

3.1 Nonlinear temporal data generation

Our primary goal is to generate nonlinear temporal datasets, which are acquired by solving associated partial differential equation (PDE) problems. Numerical methods for PDE solving are categorized into explicit and implicit schemes; these methods discretize PDE problems by mapping them from infinite-dimensional function spaces to finite-dimensional spaces, ultimately resulting in either a system of linear equations or a sequence of matrices. To illustrate this, we take the Navier-Stokes equations 1 and the Crank-Nicolson method—a classic implicit scheme—as an example.

$$u_t + \mathbf{v} \cdot \nabla u = \nu \nabla^2 u + f,\tag{1}$$

Generating datasets for temporal PDEs typically relies on the Crank-Nicolson (CN) scheme—an implicit temporal discretization method ideal for balancing stability and efficiency. Its core principle is to discretize the time derivative using a weighted average of solutions at two consecutive time steps $(t_n \text{ and } t_{n+1})$, where the time step $\Delta t = t_{n+1} - t_n$. For a temporal PDE with linear term $Lu = \nu \nabla^2 u$ (where ∇^2 denotes the Laplacian) and nonlinear term $N(u) = -\mathbf{v} \cdot \nabla u + f$, the CN scheme yields:

$$\frac{u^{n+1} - u^n}{\Delta t} = \frac{1}{2} \left(Lu^{n+1} + N(u^{n+1}) \right) + \frac{1}{2} \left(Lu^n + N(u^n) \right), \tag{2}$$

Since the nonlinear term at the t_{n+1} step is difficult to handle directly, a semi-implicit approach is adopted: $N(u^{n+1})$ is replaced with $N(u^n)$. Rearranging the equation ultimately gives:

$$u^{n+1} = \frac{\left(1 + \frac{L}{2}\Delta t\right)u^n + N(u^n)}{1 - \frac{L}{2}\Delta t},\tag{3}$$

Here, u^n and u^{n+1} are solutions at t_n and t_{n+1} ; the nonlinear term N(u) uses u^n (known from the previous step) to avoid complex coupling. Rearranging and solving the equation (often in the frequency domain, where ∇^2 becomes lap) gives u^{n+1} . To build the dataset, start from the initial

condition u^0 , iterate this process for T time steps, and collect the spatiotemporal solution pairs $\{(t_n, u^n, f^n)\}_{n=0}^T$. The time complexity of dataset generation with CN is dominated by computing multiple matrix multiplications at each time step. Let n = dimension of the spatial discretization, T = total time steps, and N = number of data points. For sparse matrices, each computation solve costs $O(Nn^2)$. Thus, the total complexity is $O(TNn^2)$.

3.2 Nonliear temporal datasets

Nonlinear temporal partial differential equations (PDEs) are fundamental mathematical tools for describing complex spatiotemporal dynamic systems across science and engineering (e.g., fluid flow, heat transfer, and chemical reactions). Their general form is typically expressed as:

$$\frac{\partial u(\boldsymbol{x},t)}{\partial t} = \mathcal{L}(u(\boldsymbol{x},t)) + \mathcal{N}(u(\boldsymbol{x},t)) + f(\boldsymbol{x},t) \qquad \qquad \boldsymbol{x} \in \Omega, \tag{4}$$

$$u(\boldsymbol{x},t) = g(\boldsymbol{x},t) \qquad \qquad \boldsymbol{x} \in \partial\Omega, \tag{5}$$

$$\nabla u(\boldsymbol{x},t) \cdot \mathbf{n} = h(\boldsymbol{x},t) \qquad \qquad \boldsymbol{x} \in \partial\Omega, \tag{6}$$

where: $\boldsymbol{x}=(x_1,x_2,...,x_d)\in\Omega$ denotes the d-dimensional spatial coordinate (with Ω as the spatial domain, e.g., $\Omega=[0,1]^2$ for 2D problems), $t\in[0,T]$ is the time variable (with T as the total time horizon), and $u(\boldsymbol{x},t)$ is the unknown solution field (e.g., velocity for fluid flow, temperature for heat transfer) that varies with both space and time. $\mathcal{L}(\cdot)$ represents the linear spatial operator (e.g., $\mathcal{L}(u)=\nu\nabla^2 u$, where ν is the diffusion coefficient and ∇^2 is the Laplacian operator describing linear diffusion). $\mathcal{N}(\cdot)$ is the nonlinear operator that introduces complexity (e.g., $\mathcal{N}(u)=-\mathbf{u}\cdot\nabla u$ for convection in Navier-Stokes equations, where the product of the solution \mathbf{u} and its spatial gradient ∇u leads to nonlinearity). Finally, $f(\boldsymbol{x},t)$ is the external source/sink term (e.g., a heat source in thermal PDEs) that drives or modifies the system's evolution. To ensure the uniqueness of the solution, boundary conditions are imposed on the domain boundary $\partial\Omega$: common types include Dirichlet conditions (specifying $u(\boldsymbol{x},t)=g(\mathbf{x},t)$ for $\boldsymbol{x}\in\partial\Omega$, where g is a given function) and Neumann conditions (specifying $\nabla u(\boldsymbol{x},t)\cdot\mathbf{n}=h(\boldsymbol{x},t)$ for $\boldsymbol{x}\in\partial\Omega$, where \mathbf{n} is the outward unit normal vector of $\partial\Omega$ and h is a given function).

4 METHOD

In existing algorithms, we typically generate initial conditions and RHS of nonlinear temporal PDEs randomly, then feed them into traditional solvers for such PDEs to obtain the solution function u(x). However, taking the semi-implicit method as mentioned in 3.1 as an example, we have to iterate over thousands of time steps to ensure solution stability —this demands significant computational resources and time. Yet in practice, datasets used to train neural operators typically only require the solution at a few to a dozen time steps. Thus, a natural question emerges: can we reduce the number of iterations to accelerate dataset generation?

Thus, unlike traditional methods, HOPSS acts directly on training-level data, thereby avoiding iteration over a large number of time steps. The specific process is as follows: we first generate a set of solutions using a high-precision solver, then align them to the time steps required for training via downsampling—these serve as the base solutions in our solution space. Next, the base functions are fed into the generator, where two solutions are randomly selected: one is scaled by a small scalar to act as the homologous perturbation term, and then combined with the other to generate a new solution, along with small random noise. Finally, based on these new solutions and their corresponding physical equations, we compute the right-hand side (RHS) of the equation; these collectively form new solution pairs. This strategy acts directly on training-level data, with the only exception being the generation of base solutions. Thus, when generating new data, it avoids the need to iterate over a large number of time steps.

Figure 1: Overview of the HOPSS method. First, generating a series of solutions using the existing method and downsampling as base solutions. Then, randomly select two base solutions and apply the homologous perturbation to generate new solutions. Finally, calculating the RHS using the new solutions.

The HOPSS method primarily comprises three key steps:

- 1. **Base solution generation**: In this phase, we generate high-precision solution functions using a high-precision solver and format them to meet the requirements for model training.
- 2. **Homologous perturbation in solution space**: In this stage, we apply homologous perturbation to the base solutions to generate a large number of new candidate solutions.
- 3. **Computation of equation RHS**: In this phase, we calculate the corresponding right-hand side (RHS) of the physical equation based on the generated candidate solutions, ensuring they satisfy the constraints of the target physical equation.

4.1 SOLUTION FUNCTIONS GENERATION

The HOPSS method generates a set of base functions by first producing solutions based on a designated distribution consistent with real-world physical scenarios—these base functions typically number 100 to 500, denoted herein as N_{base} . For the generation of these solutions, this paper adopts existing high-precision methods. Subsequently, we downsample the generated solutions to align them with the time steps required for training, reducing the number of time steps from several thousand to several dozen. These processed solutions serve as the base functions of HOPSS and constitute the foundational elements of its solution space.

4.2 Homologous Perturbation in Solution Space

The homologous perturbation process entails introducing small-scale functional perturbations to existing base solution functions. Specifically, we randomly select two base solution functions: one acts as the primary function u_i , and the other as the homologous perturbation term, denoted u_j . To prevent the generated new solution function from exhibiting excessive fluctuations, we multiply the perturbation term u_j by a small constant μ (typically $\mu \approx 10^{-3}$). After combining these two functions, we further add small-scale time-invariant random noise ξ —specifically Gaussian noise where $\xi \sim \mathcal{N}(0,\varsigma I)$ with $\varsigma \approx 10^{-4}$ —to enhance the robustness and diversity of the dataset. The mathematical expression for this process is:

$$u_{\text{new}} = u_i + \mu \cdot u_j + \xi. \tag{7}$$

4.3 COMPUTE VARIATION OF RHS

To satisfy the constraints of the physical equations, we recompute the variation of the right-hand side (RHS) at the training data level. Using the same discretization method as employed earlier, we derive a system of equations corresponding to the target physical equation. We then substitute

the newly generated u_{new} and the primary function u_i into this discretized system of equations—this enables us to calculate the variation Δf of the original RHS term f_i , and further obtain the new RHS term f_{new} . Subsequently, u_{new} and f_{new} form new solution pairs, which serve as the new dataset.

Taking a burgers equation with a source term as an example (see Eq. 8):

$$\frac{\partial u}{\partial t} + uu_x = \nu u_{xx} + f(x, t),\tag{8}$$

We assume both $(u_{\text{new}}, f_{\text{new}})$ and (u_i, f_i) satisfy Eq. 8. According to the generation rule for u_{new} , we introduce the transformation:

$$u_i = u_{\text{new}} - v, \tag{9}$$

where $v = \mu u_j + \xi$ is a newly defined composite variable, with μ being a constant and $\xi = \xi(x)$ being a spatially dependent function.

By substituting $u_i = u_{\text{new}} - v$ into Eq. 8 and rearranging terms, we can derive that u_{new} satisfies the same form of Burgers equation:

$$\frac{\partial u_{\text{new}}}{\partial t} + u_{\text{new}} u_{\text{new},x} = \nu u_{\text{new},xx} + f_{\text{new}}, \tag{10}$$

where the $u_{new,x} = \frac{\partial u_{new}}{\partial x}$, the $u_{new,xx} = \frac{\partial^2 u_{new}}{\partial x^2}$ and the new source term f_{new} is given by:

$$f_{\text{new}} = f_i + \underbrace{\frac{\partial v}{\partial t} + (u_{new}v)_x - vv_x - \nu v_{xx}}_{\Delta f}.$$
 (11)

5 THEORETICAL ANALYSIS

5.1 Existing Method

Solving temporal PDEs inherently involves spatiotemporal discretization, where the primary computational expense stems from iterative time step evolution and the associated solution of spatially discretized linear systems LeVeque (2007); Hughes (2000). Conventional numerical methods for temporal PDEs typically combine temporal discretization (e.g., 4th-order Runge-Kutta (RK4) Hairer et al. (1993), Crank-Nicolson (CN) scheme Crank & Nicolson) and spatial discretization (e.g., Finite Element Method (FEM) Hughes (2000), Finite Difference Method (FDM) LeVeque (2007)); the computational complexity of these methods is determined by the interplay of spatial grid scale, time step count, and linear system solving cost.

Notably, for the semi-impact scheme, the cost of time iterations dominates the overall computational process: for the traditional method, with n as the space dimensions, this step exhibits a time complexity of $O(n^2)$ Dong et al. (2024), and a linear system must be solved at each time step to update the PDE solution. Given T total time steps, usually exceeding thousands, and n as the dimension of the spatially discretized linear system, the overall time complexity of the entire solving process amounts to $O(Tn^2)$ for dense matrix scenarios. If N datapoints are generated, the all-time complexity comes up to $O(NTn^2)$.

5.2 OUR HOPSS METHOD

According to the introduction referred to in 4, our method consists of three steps: base solution generation, variational action, and RHS computation.

In the first step, we use traditional methods to generate N_b base solutions. According to the forward analysis, its time complexity is $O(N_bTn^2)$. Since the cost of the downsampling step is far lower than that of data generation, we reasonably ignore it. We then operate on training-level data, which

typically has smaller T' (time steps) and n' (spatial dimensions) than the T and n used in traditional data generation. Typically, T' is $\frac{T}{1000}$ or smaller.

In the subsequent two steps, the time complexity is $O(N_{\text{new}}T'(n')^2)$ obviously. Since T' is much smaller than T and can be approximated as a constant, the time complexity for subsequent new data generation is approximately $O(N_{\text{new}}(n')^2)$ —which is one order of magnitude lower than that of traditional methods.

Thus, the total time complexity of the HOPSS method is $O(N_bTn^2 + N_{\text{new}}(n')^2) \approx O(N_bTn^2)$. Consequently, the theoretical acceleration ratio of HOPSS depends on the gap between N_b (the number of base solutions) and N (the number of samples required for generating an equivalent dataset via traditional methods).

6 EXPERIMENT

In this chapter, we compare our proposed data generation method with existing data generation methods.

6.1 Experiment setup

Our analysis focuses on two key performance indicators, both critical for evaluating the effectiveness of data generation methods:

- Time cost of data generation
- Test loss of neural operator models trained on the generated data

In our experiments, we test two widely recognized and adopted neural operator models—among the most prominent and prevalent in data-driven PDE algorithms:

- FNO (Fourier Neural OperatorLi et al. (2021))
- Transolver(Wu et al. (2024))

We also evaluate three types of PDE problems with significant applications in science and engineering:

- Navier-Stokes equations(Li et al. (2021))
- Burgers' equation(Xie et al. (2013))
- KdV equation (Korteweg-de Vries equationShen (1993))

Baselines. The primary time cost of existing data generation methods stems from iterating over thousands of time steps, despite only dozens of time steps being required for model training. We use the traditional methods to generate solution functions, which serve as our baselines. The details of the experiment are shown in AppendixB.1 and parameters of the generated dataset in AppendixB.3.

6.2 Main result

Our main results across all datasets and models are presented in Table 1. Further details and hyperparameters are provided in the AppendixB.2. Based on these results, we draw the following observations.

Firstly, the HOPSS method exhibits significant acceleration compared to traditional methods—particularly for the Navier-Stokes equations, where the acceleration ratio reaches 10 times. The time cost of our method consists of three components: generating basis functions, performing homologous perturbation, and computing the RHS. Generating basis functions requires iterating over thousands of time steps, which are then downsampled to dozens of time steps. The other stages involve operating on data with dozens of time steps. Thus, the primary time cost of the entire process lies in generating the basis functions. This implies that our HOPSS method can generate large volumes of training data at low cost, underscoring the efficiency of our approach.

Table 1: Performance comparison between traditional methods and HOPSS on different PDE problems. The first column lists the method used to generate the dataset and the number of training instances. The first row represents the corresponding PDE problem and the training models.

Method	N	avier-St	okes	KdV			Burgers		
	TIME(s)	FNO	Transolver	TIME(s)	FNO	Transolver	TIME(s)	FNO	Transolver
Tradition1000	1.20e4	6.7e - 3	$1.4e{-2}$	2.16e3	8.7e - 3	$3.1e{-2}$	4.29e4	6.1e-2	5.9e-2
HOPSS1000	1.22e3	7.6e - 3	$1.5\mathrm{e}{-2}$	1.08e3	1.7e-2	$3.6e{-2}$	2.15e4	6.2e-2	7.2e-2
HOPSS10000	1.38e3	3.3e - 3	$8.8e{-3}$	1.09e3	9.8e - 3	$3.8e{-2}$	$2.16\mathrm{e}4$	4.7e - 2	$7.3e{-2}$

Secondly, across different PDE problems and with various neural operators, datasets generated by the HOPSS method exhibit comparable performance to those from traditional methods. For example, in the case of the Navier-Stokes equations, Transolver and FNO achieve nearly identical performance when trained on datasets of the same size. Moreover, as dataset sizes increase, our method can even outperform traditional methods. This demonstrates that the HOPSS method simultaneously accelerates dataset generation while maintaining approximate precision.

6.3 HYPERPARAMENTS ANALYSIS

Table 2: Burgers equation hyperparameter test results. Left: Influence of perturbation level μ (with fixed Gaussian noise and $N_b=100$). Middle: Influence of the number of base solutions N_b (with fixed Gaussian noise and $\mu=10^{-3}$). Right: Influence of noise type (with fixed $\mu=10^{-3}$ and $N_b=500$). Test loss is obtained from FNO.

Perturbation Level	Test Loss	$N_{\rm b}$	Test Loss	-	Noise Type	Test Loss
10^{-1}	0.182	100	0.0783		Gaussian	0.0621
10^{-2}	0.0787	200	0.0652		Perlin	0.0625
10^{-3}	0.0783	300	0.0650		Multi_sine	0.0624
10^{-4}	0.0784	400	0.0639		Random_walk	0.0622
10^{-5}	0.0790	500	0.0621		Zero	0.0628

This section will show how to influence the dataset performance for the different hyperparameters.

Perturbation level μ : As illustrated in the first subtable, dataset performance degrades as the perturbation level increases. In this experiment, we fixed the noise type (Gaussian noise) and used 100 basis solutions for both the Burgers equation and the FNO model. When the perturbation level is set to 0.1, the test loss reaches 0.182, indicating significant performance degradation. In contrast, reducing the perturbation level to 0.001 lowers the test loss to 0.0783, a substantial improvement. This suggests that excessively large perturbations may cause the generated data to become more likely out-of-distribution, thereby undermining the model's ability to learn meaningful patterns.

Number of base solutions N_b : Consistent with the results in the second subtable, dataset quality exhibits a strong correlation with the number of base solutions. In this experiment, we fixed the noise type (Gaussian noise) and set the perturbation level at 0.001—with tests conducted on both the Burgers equation and the FNO model. As the number of base solutions increases from 100 to 500, the test loss decreases from 0.0783 to 0.0621, resulting in a performance improvement of approximately 20%. This trend indicates that a larger number of base solutions helps construct a more complete subspace of the solution space, thereby enhancing the quality of the dataset.

Noise type: From the third subtable, it is evident that noise type has a negligible impact on model performance. The details of the noise are shown in AppendixD. In this experiment, we fixed the number of basis solutions at 500 and set the perturbation level to 0.001—consistent across both the Burgers equation setup and the FNO model. We evaluated five experimental scenarios: four distinct noise types (Gaussian, Perlin, Multi-sine, and Random-walk noise) and a no-noise condition, and the resulting test loss ranged from 0.621 to 0.628, no more than 5%. This consistency underscores that our method demonstrates strong robustness to variations in noise type.

6.4 TIME STEPS INFLUENCE ON TRADITIONAL METHOD

In this section, we analyze the influence of different time steps on the traditional method, focusing specifically on two key metrics: dataset generation time cost and the performance of the subsequently trained models. As shown in Table 3, time steps exert a significant impact on both generation time cost and model training performance. For time cost: as the time step decreases from 5×10^{-3} to 1×10^{-3} , the generation time cost increases nearly 9-fold from 1317.47 s to 11980.42 s. In terms of model performance, the FNO test loss decreases substantially from 0.106 to 0.0067, representing a marked im-

Table 3: Performance comparison under different time step sizes for the Navier-Stokes problem in the traditional method

Time Step Sizes	Time Cost (s)	Test Loss
1×10^{-3}	11980.42	0.0067
2×10^{-3}	4380.12	0.024
5×10^{-3}	1317.47	0.106

provement in prediction accuracy. This observation reveals a clear trade-off: smaller time steps yield better model performance but come at the cost of significantly higher dataset generation time.

6.5 ABLATION RESULT

 Finally, we performed an ablation study to illustrate the critical impact of the solution generation component in our HOPSS method on dataset quality. As shown in Table 4, we evaluated datasets generated via the Mixup method and our HOPSS method (for detailed implementation of Mixup, refer to Appendix C). Experimental results confirm the effectiveness of HOPSS in high-quality dataset generation: when models are trained on Mixup-generated data, the test errors for FNO and Transolver reach 0.312 and 1.53, respectively—values that indicate

Table 4: Ablation experiment results comparing datasets generated using different methods

Method	FNO	Transolver
Mixup	0.312	1.53
HOPSS	0.0076	0.015

the Mixup-generated dataset offers minimal training value. In contrast, our HOPSS method proves effective in generating datasets efficiently while ensuring accurate model predictions.

7 CONCLUSION AND FUTURE WORK

Conclusion: In this paper, we propose the HOPSS method for generating datasets tailored to nonlinear temporal PDEs. Specifically, the method comprises three key steps: base solution generation, homologous perturbation in the solution space, and computation of the equation's right-hand side. By acting directly on the dozens of time steps required for training and computing the corresponding RHS, the HOPSS method accelerates the data generation process while preserving the precision of the data used for model training. This approach effectively overcomes a major barrier in dataset generation for nonlinear temporal PDEs.

Future Work: Building on HOPSS's promising performance in accelerating nonlinear temporal PDE dataset generation, we will advance the work via two targeted directions: First, we will optimize base solution selection to enhance solution space quality. Specifically, we will introduce solution space coverage analysis and active learning: by quantifying initial base solution representativeness, we will develop a framework to filter optimal ones. This prioritizes filling coverage gaps, boosting base solution diversity, and physical comprehensiveness to enhance perturbation and dataset generalization. Second, we will develop a unified physics-informed metric for dataset quality, moving beyond over-reliance on indirect indicators like model test loss. This integrated metric will encompass dual dimensions: physical consistency (e.g., adherence to PDE conservation laws, error against high-precision references) and data utility (e.g., sample diversity, cross-PDE generalization capacity). These will be integrated into a unified pipeline for comprehensive dataset assessment to guide HOPSS optimization.

ETHICS STATEMENT

We have manually reevaluated the dataset we created to ensure it is free of any potential for discrimination, human rights violations, bias, exploitation, and any other ethical concerns.

REPRODUCIBILITY STATEMENT

To ensure the reproducibility of our findings, all source code and datasets used in our experiments are included in the supplementary material. The provided materials are sufficient to replicate the main results presented in this paper.

REFERENCES

- Johannes Brandstetter, Max Welling, and Daniel E. Worrall. Lie point symmetry data augmentation for neural pde solvers, 2022.
- J. Crank and P. Nicolson. A practical method for numerical evaluation of solutions of partial differential equations of the heat-conduction type. 6(1):207–226. ISSN 1572-9044. doi: 10.1007/BF02127704. URL https://doi.org/10.1007/BF02127704.
- Huanshuo Dong, Hong Wang, Haoyang Liu, Jian Luo, and Jie Wang. Accelerating pde data generation via differential operator action in solution space. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24. JMLR.org, 2024.
- Sergei K Godunov and Ihor Bohachevsky. Finite difference method for numerical computation of discontinuous solutions of the equations of fluid dynamics. *Matematičeskij sbornik*, 47(3):271–306, 1959.
- Ernst Hairer, Syvert Norsett, and Gerhard Wanner. Solving Ordinary Differential Equations I: Non-stiff Problems, volume 8. 01 1993. ISBN 978-3-540-56670-0. doi: 10.1007/978-3-540-78862-1.
- Y. Hao, Z. Li, and G.E. Karniadakis. Data-efficient learning for pde solvers via physics-informed redundancy reduction. *Journal of Computational Physics*, 462:111264, 2022. doi: 10.1016/j.jcp. 2022.111264.
- Jun-Ting Hsieh, Shengjia Zhao, Stephan Eismann, Lucia Mirabella, and Stefano Ermon. Learning neural pde solvers with convergence guarantees, 2019. URL https://arxiv.org/abs/1906.01200.
- Thomas JR Hughes. *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*. Dover Publications, 2000. ISBN 9780486411811.
- 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), May 2021. ISSN 1091-6490. doi: 10.1073/pnas.2101784118. URL http://dx.doi.org/10.1073/pnas.2101784118.
- Randall J LeVeque. *Finite Difference Methods for Ordinary and Partial Differential Equations: Steady-State and Time-Dependent Problems.* Society for Industrial and Applied Mathematics, 2007. doi: 10.1137/1.9780898717839.
- R.J. LeVeque. *Finite Volume Methods for Hyperbolic Problems*. Cambridge University Press, 2002. doi: 10.1017/CBO9780511791253.
- 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*, 2021. URL https://openreview.net/forum?id=c8P9NQVtmnO.

- Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, March 2021. ISSN 2522-5839. doi: 10.1038/s42256-021-00302-5. URL http://dx.doi.org/10.1038/s42256-021-00302-5.
 - K.W. Morton and D.F. Mayers. *Numerical Solution of Partial Differential Equations: An Introduction*. Cambridge University Press, 2005. doi: 10.1017/CBO9780511619688.
 - Samuel S Shen. Forced kdv equation. In *A Course on Nonlinear Waves*, pp. 147–187. Springer, 1993.
 - Kimberly L. Stachenfeld, Drummond B. Fielding, Dmitrii Kochkov, Miles Cranmer, Tobias Pfaff, Jonathan Godwin, Can Cui, Shi-Lin Ho, Peter W. Battaglia, and Alvaro Sanchez-Gonzalez. Learned coarse models for efficient turbulence simulation. *ArXiv*, abs/2112.15275, 2021. URL https://api.semanticscholar.org/CorpusID:245634354.
 - Gilbert Strang, George J Fix, et al. *An analysis of the finite element method*, volume 212. Prenticehall, 1973.
 - J.C. Strikwerda. Finite Difference Schemes and Partial Differential Equations. Society for Industrial and Applied Mathematics (SIAM), 2nd edition, 2004. doi: 10.1137/1.9780898717839.
 - Roger Temam. *Navier–Stokes equations: theory and numerical analysis*, volume 343. American Mathematical Society, 2024.
 - Hong Wang, Zhongkai Hao, Jie Wang, Zijie Geng, Zhen Wang, Bin Li, and Feng Wu. Accelerating data generation for neural operators via krylov subspace recycling, 2024. URL https://arxiv.org/abs/2401.09516.
 - Haixu Wu, Huakun Luo, Haowen Wang, Jianmin Wang, and Mingsheng Long. Transolver: A fast transformer solver for pdes on general geometries, 2024. URL https://arxiv.org/abs/2402.02366.
 - Huantian Xie, Dingfang Li, and Feng Li. A new numerical method of particular solutions for inhomogeneous burgers' equation. *Mathematical Problems in Engineering*, 2013(1):974808, 2013.
 - Cheng Yang, Xubo Yang, and Xiangyun Xiao. Data-driven projection method in fluid simulation. *Computer Animation and Virtual Worlds*, 27(3-4):415–424, 2016. doi: https://doi.org/10.1002/cav.1695. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/cav.1695.
 - Xuan Zhang, Limei Wang, Jacob Helwig, Youzhi Luo, Cong Fu, Yaochen Xie, Meng Liu, Yuchao Lin, Zhao Xu, Keqiang Yan, Keir Adams, Maurice Weiler, Xiner Li, Tianfan Fu, Yucheng Wang, Alex Strasser, Haiyang Yu, YuQing Xie, Xiang Fu, Shenglong Xu, Yi Liu, Yuanqi Du, Alexandra Saxton, Hongyi Ling, Hannah Lawrence, Hannes Stärk, Shurui Gui, Carl Edwards, Nicholas Gao, Adriana Ladera, Tailin Wu, Elyssa F. Hofgard, Aria Mansouri Tehrani, Rui Wang, Ameya Daigavane, Montgomery Bohde, Jerry Kurtin, Qian Huang, Tuong Phung, Minkai Xu, Chaitanya K. Joshi, Simon V. Mathis, Kamyar Azizzadenesheli, Ada Fang, Alán Aspuru-Guzik, Erik Bekkers, Michael Bronstein, Marinka Zitnik, Anima Anandkumar, Stefano Ermon, Pietro Liò, Rose Yu, Stephan Günnemann, Jure Leskovec, Heng Ji, Jimeng Sun, Regina Barzilay, Tommi Jaakkola, Connor W. Coley, Xiaoning Qian, Xiaofeng Qian, Tess Smidt, and Shuiwang Ji. Artificial intelligence for science in quantum, atomistic, and continuum systems. Foundations and Trends® in Machine Learning, 18(4):385–912, 2025. ISSN 1935-8245. doi: 10.1561/2200000115. URL http://dx.doi.org/10.1561/2200000115.

A USAGE OF LLMS

Throughout the preparation of this manuscript, Large Language Models (LLMs) were utilized as a writing and editing tool. Specifically, we employed LLMs to improve the clarity and readability of the text, refine sentence structures, and correct grammatical errors. All final content, including the core scientific claims, experimental design, and conclusions, was conceived and written by us, and we take full responsibility for the final version of this paper.

B SPECIFIC EXPERIMENTAL DETAILS

B.1 HARDWARE SETUP

The data generation is running on Intel(R) Xeon(R) Silver 4316, and the models are training on a GeForce RTX 4090 GPU with 24GB of memory.

B.2 MODEL SET

FNO1d: we employ 3 FNO layers with learning rate 0.0001, batch size 20, epochs 1000, modes 16, and width 20.

FNO2d: we employ 4 FNO layers with learning rate 0.001, batch size 60, epochs 500, modes30, and width 60.

Transolver_Irregular_Mesh: we set 3 hidden layers with 64 dimensions for each hidden layer, batch size 4, heads 8, slice nums 32, learning rate 0.002, and epochs 1000.

Transolver_Structured_Mesh_2D: we set 3 hidden layers with 64 dimensions for each hidden layer, batch size 8, heads 4, slice nums 32, learning rate 0.001, and epochs 500.

B.3 DATA

B.3.1 NAVIER-STOKES EQUATIONS

In this research, we delve into two-dimensional Navier-Stokes equations for a viscous, incompressible fluid in vorticity form on the unit torus, which are governed by the equation(Li et al. (2021)).

$$\begin{cases} \partial_t w(x,t) + u(x,t) \cdot \nabla w(x,t) = \nu \Delta w(x,t) + f(x,t), & x \in (0,1)^2, \ t \in (0,T] \\ \nabla \cdot u(x,t) = 0, & x \in (0,1)^2, \ t \in [0,T] \\ w(x,0) = w_0(x), & x \in (0,1)^2 \end{cases}$$
(12)

where $u \in C([0,T];H^r_{\mathrm{per}}((0,1)^2;\mathbb{R}^2))$ for any r>0 is the velocity field, $w=\nabla\times u$ is the vorticity, $w_0\in L^2_{\mathrm{per}}((0,1)^2;\mathbb{R})$ is the initial vorticity, $\nu\in\mathbb{R}_+$ is the viscosity coefficient, and $f\in L^2_{\mathrm{per}}((0,1)^2;\mathbb{R})$ is the forcing function. For our experimental setup, the dataset is generated by the Crank-Nicolson method. We generate the dataset with $gridsize=128, T=10, \nu=1e^{-4}$ and $\Delta t=1e^{-3}$. Then we downsample the dataset to the gridsize=64 and $\Delta t=0.5$ for training the model. The force function is generated from a Gaussian Random Field (GRF) methodology, with a time constant $\tau=2.0$ and a decay exponent $\alpha=2.5$. In the HOPSS method, we use 100 solution functions as basis functions.

B.3.2 Burgers Equations

In this research, we investigate one-dimensional Burgers equations, expressed as (Xie et al. (2013)):

$$\frac{\partial u(x,t)}{\partial t} + u(x,t)\frac{\partial u(x,t)}{\partial x} - \frac{1}{R}\frac{\partial^2 u(x,t)}{\partial x^2} = f(x,t). \tag{13}$$

Here, u(x,t) denotes the velocity field; R is the Reynolds number, a dimensionless parameter characterizing the ratio of inertial forces to viscous forces; and f(x,t) represents the time-dependent forcing function. For our experimental setup, R is set to 1000, indicating a flow regime with relatively strong inertial effects, while the forcing function f(x,t) is generated using a Gaussian

Random Field (GRF) method with key parameters: $\tau=7$ (time scale, controlling the temporal correlation of the random field), $\alpha=2.5$ (decay exponent, governing the spatial correlation decay rate), and $\sigma=7^2$ (variance, determining the amplitude of random fluctuations). For numerical discretization, the original spatial grid size (number of spatial sampling points) is 1024 with a time step $\Delta t=5\times10^{-3}$ (temporal resolution), which we downsample to a grid size of 64 and $\Delta t=5\times10^{-2}$ for model training, and in the HOPSS method, 500 solution functions are adopted as basis functions.

B.3.3 Forced KdV Equations

 We also study one-dimensional forced Korteweg-de-Vries (KdV) equations, expressed as (Shen (1993)):

$$u_t + \lambda u_x + 2\alpha u u_x + \beta u_{xxx} = f'(x), \quad -\infty < x < \infty$$
 (14)

Where u(x,t) is the wave amplitude; λ is the linear advection coefficient; α is the nonlinear coefficient (governing wave steepening); β is the dispersion coefficient (counteracting steepening via wave dispersion); and f'(x) denotes the spatial derivative of the forcing function f(x). For our experimental setup, coefficients are set as $\alpha=-0.5$, $\beta=-1.0$, and $\lambda=0$, configuring the equation to describe weakly nonlinear dispersive waves, while f'(x) is directly generated via GRF with parameters: $\tau=5.0$ (time scale), $\alpha=2.5$ (spatial decay exponent), and $\sigma=1.0$ (variance of the random field). For numerical discretization, the original spatial grid size is 512 with 10000 time steps (total temporal sampling points), which are downsampled to a grid size of 64 and 20 time steps for model training, and in the HOPSS method, 500 solution functions are used as basis functions.

C MIXUP METHOD

The mathematical equation can be expressed as:

$$u_{new} = \sum_{i=0}^{N_b} \alpha_i u_i, \sum_{i=0}^{N_b} \alpha_i = 1$$
 (15)

where u_i is the solutions and α_i is a constant. In our experiment, α_i is sampled from a standard Gaussian distribution and then normalized. N_b is set to 100.

D NOISY INTRODUCTION

We corrupt clean data x (a vector or the last dimension of a tensor) with several synthetic noise models. A unified relative amplitude parameter $\varepsilon = \texttt{noise_level}$ is mapped to an absolute scale

$$A = \varepsilon \cdot \max_{i} |x_i|, \qquad A > 0 \text{ (fallback } A = 10^{-8} \text{ if } \max |x_i| = 0). \tag{16}$$

All raw noise patterns $\tilde{\eta}$ are rescaled to $\eta = A \tilde{\eta} / \max |\tilde{\eta}|$ (except the Gaussian case which is sampled directly with variance A^2). The noisy signal is $x' = x + \eta$.

D.1 GAUSSIAN NOISE

Standard i.i.d. zero-mean Gaussian (baseline):

$$\eta_i^{(G)} \sim \mathcal{N}(0, A^2), \quad x_i' = x_i + \eta_i^{(G)}.$$
 (17)

Parameters used: only ε (e.g. $\varepsilon = 10^{-3}$). No spatial correlation; flat spectrum.

D.2 MULTI-SINE NOISE

A smooth low-frequency superposition of K sinusoidal modes with random coefficients and phases over normalized coordinate $s \in [0, 1]$ discretized into L points:

$$\tilde{\eta}(s) = \sum_{k=1}^{K} \left[a_k \sin(2\pi k s + \phi_k) + b_k \cos(2\pi k s + \psi_k) \right],$$
(18)

With $a_k, b_k \sim \mathcal{U}(-1, 1)$ and phases ϕ_k, ψ_k sampled uniformly (implemented as a single phase per mode). Normalization:

$$\eta^{(MS)}(s) = A \frac{\tilde{\eta}(s)}{\max_s |\tilde{\eta}(s)|}.$$

Parameters: K = multisine = 8; relative amplitude ε . Result: band-limited (low-frequency) non-Gaussian smooth perturbation.

D.3 PERLIN NOISE

One-dimensional Perlin procedural noise with C lattice cells (C+1 gradient points). Let $s \in [0,1]$, t = sC, $i = \lfloor t \rfloor$, $u = t - i \in [0,1)$. Random gradients $G_i \sim \mathcal{U}(-1,1)$. Use the quintic fade

$$fade(u) = 6u^5 - 15u^4 + 10u^3.$$

Define endpoint contributions $v_0 = G_i u$, $v_1 = G_{i+1}(u-1)$ and interpolate:

$$\tilde{\eta}(s) = v_0 + (v_1 - v_0) \operatorname{fade}(u), \qquad \eta^{(P)}(s) = A \frac{\tilde{\eta}(s)}{\max_s |\tilde{\eta}(s)|}.$$

Parameters: $C=\min(\text{perlin_cells},L-1)$ with $\text{perlin_cells}=32$; relative amplitude ε . Produces locally smooth, multi-scale, non-Gaussian structure.

D.4 RANDOM WALK NOISE

Cumulative uniform increments (strong correlation, non-stationary before centering):

$$\delta_{\ell} \sim \mathcal{U}(-1,1), \quad \tilde{\eta}_{\ell} = \sum_{j=1}^{\ell} \delta_{j}, \quad \bar{\eta} = \frac{1}{L} \sum_{\ell=1}^{L} \tilde{\eta}_{\ell}, \quad \hat{\eta}_{\ell} = \tilde{\eta}_{\ell} - \bar{\eta}, \quad \eta_{\ell}^{(\mathrm{RW})} = A \frac{\hat{\eta}_{\ell}}{\max_{\ell} |\hat{\eta}_{\ell}|}.$$

Parameters: only ε (no extra hyperparameter). Produces low-frequency drift-like perturbations after mean removal.

Broadcast over higher dimensions. For tensors shaped (N, S) or (N, S, T), the non-Gaussian patterns are generated along the last axis (S) and broadcast to other leading dimensions; per-sample independent noise can be obtained by generating a separate pattern per batch slice.