
FMINT: BRIDGING HUMAN DESIGNED AND DATA PRETRAINED MODELS FOR DIFFERENTIAL EQUATION FOUNDATION MODEL

Zezheng Song* **Jiaxin Yuan***
 Department of Mathematics
 University of Maryland College Park
 College Park, MD, USA
 {zsong001, jyuan98}@umd.edu

Haizhao Yang
 Department of Mathematics
 Department of Computer Science
 University of Maryland College Park
 College Park, MD, USA
 hzyang@umd.edu

ABSTRACT

Human-designed algorithms have long been fundamental in solving a variety of scientific and engineering challenges. Recently, data-driven deep learning methods have also risen to prominence, offering innovative solutions across numerous scientific fields. While traditional algorithms excel in capturing the core aspects of specific problems, they often lack the flexibility needed for varying problem conditions due to the absence of specific data. Conversely, while data-driven approaches utilize vast datasets, they frequently fall short in domain-specific knowledge. To bridge these gaps, we introduce **FMint** (Foundation Model based on Initialization), a generative pre-trained model that synergizes the precision of human-designed algorithms with the adaptability of data-driven methods. This model is specifically engineered for high-accuracy simulation of dynamical systems. Starting from initial trajectories provided by conventional methods, FMint quickly delivers highly accurate solutions. It incorporates in-context learning and has been pre-trained on a diverse corpus of 500,000 dynamical systems, showcasing exceptional generalization across a broad spectrum of real-world applications. By effectively combining algorithmic rigor with data-driven flexibility, FMint sets the stage for the next generation of scientific foundation models, tackling complex problems with both efficiency and high accuracy.

Keywords Foundation Model · Dynamical Systems · Fast Simulation · In-context Learning

1 Introduction

Throughout history, humanity has developed a myriad of algorithms to address real-world challenges. For instance, numerical integration methods such as the Euler and Runge–Kutta methods have significantly advanced our ability to solve ordinary differential equations (ODEs) with guaranteed accuracy. More recently, deep learning has demonstrated remarkable success in data-driven approaches across various scientific domains, including solving partial differential equations (PDEs), learning operators, and addressing inverse problems. This success raises an intriguing question: is it possible to integrate the strengths of traditional human-designed algorithms with modern data-driven methods?

Human-designed algorithms excel in capturing the structural nuances of specific problems but often lack the adaptability needed for varying parameters and conditions. Conversely, data-driven approaches effectively utilize large data sets but typically underperform in data-scarce environments and may lack essential domain knowledge. To address these challenges, we introduce **FMint**, a generative pre-trained foundation model designed to synergize the strengths of both approaches. FMint is equipped with in-context learning capabilities and has been pre-trained on over 500,000 dynamical systems, enabling it to generate highly accurate simulations from initial user-provided data. Moreover, our model delivers rapid, high-accuracy solutions leveraging its generalization capabilities and the ability to perform direct inference without the need for parameter updates. FMint represents the first foundation model to seamlessly integrate human expertise with adaptable learning, making it a pioneering tool in the simulation of diverse dynamical systems.

*Equal contribution.

Problem set-up Dynamical systems are fundamental in describing the evolution of physical states across a wide range of disciplines, including physics [1, 2, 3], chemistry [4, 5], engineering [6, 7, 8], and finance [9, 10]. Typically, these systems are formulated as systems of ODEs:

$$\frac{d\mathbf{u}(t)}{dt} = \mathbf{f}[\mathbf{u}(t)], \quad \mathbf{u}(0) = \mathbf{c}_0, \quad (1)$$

where \mathbf{c}_0 denotes the initial condition of the system. To solve this system numerically, one employs a numerical integration method given as

$$\mathbf{u}_{n+1} = \mathbf{u}_n + S(\mathbf{f}, \mathbf{u}_n, \Delta t_n), \quad \mathbf{u}_0 = \mathbf{c}_0, \quad n = 0, 1, \dots \quad (2)$$

where S represents the numerical integration scheme, such as the Euler method, and Δt_n represents the step size at n -th time step, and $\mathbf{u}_n \in \mathbb{R}^n$ is the approximated solution at cumulative time $\sum_{i=0}^n \Delta t_i$. To obtain higher simulation accuracy, users may adjust the integration scheme S and the time step Δt_n based on specific requirements. However, large-scale simulation often necessitate simulating numerous trajectories, each with unique initial conditions. In these scenarios, it is practical to standardize the time step $\Delta t := \Delta t_1 = \Delta t_2 = \dots$ across simulations, facilitating batch processing. Yet, this standardization introduces a trade-off between accuracy and efficiency: A larger time step, e.g. $k\Delta t$ speeds up the simulation at the cost of increased error, and vice versa. The trade-off underlines a key limitation of human-designed algorithms: their inability to dynamically adapt to the diverse conditions inherent in different types of problems. Although human-designed numerical integration schemes such as the Euler method and Runge–Kutta method are essential for simulating a wide array of ODEs, these methods do not fully leverage data or incorporate problem-specific information, which can restrict their effectiveness.

To address these challenges, we introduce FMInt, a novel foundation model embedded within a meta-learning framework designed for the rapid simulation of dynamical systems. This model significantly refines coarse solutions, delivering highly accurate results through fast inference. Pre-trained on a diverse dataset of 500,000 ODEs, FMInt demonstrates robust generalization across various dynamical systems, consistently achieving high-accuracy solutions.

2 Related Work

2.1 Neural network based PDE solver

In recent years, neural network methods have increasingly been applied to tackle scientific problems, including solving PDEs, operator learning, and inverse problems. For instance, Physics-Informed Neural Networks (PINNs) [11, 12, 13, 14, 15] parameterize the PDE solution within a neural network framework. The network’s loss function is constructed by enforcing the PDE constraints, integrating physical laws directly into the learning process. Similarly, the Finite Expression Method (FEX) [16, 17, 18] parameterizes solutions to PDEs using computer algebra, effectively capturing the solution structure and thus offering high accuracy and interpretable results. An alternative approach involves learning a mapping from initial conditions or varying parameters to solutions, known as the neural operator. This method achieves discretization invariance by learning a family of parametric maps, allowing it to generalize across different parameters and conditions of PDEs [19, 20, 21, 22]. Another approach involves using a neural network as an error corrector to enhance simulations. For instance, NeurVec [23] employs this strategy to enable rapid simulation of ordinary differential equations (ODEs) with large time steps. This approach achieved decent accuracy with relatively coarse integration steps on several classic dynamical systems. However, NeurVec requires specific training for each distinct ODE system using a fixed coarse step $k\Delta t$, making its retraining both costly and computationally inefficient. Moreover, its lack of generalization capability to out-of-distribution (OOD) systems significantly restricts its practicality for large-scale real-world simulations.

2.2 Foundation Model in Scientific Machine Learning

Inspired by the breakthroughs achieved by large language models like GPT-4 [24], DALL-E [25], and Llama [26], the scientific machine learning community has experienced a marked increase in the adoption of foundation models over the past year. These models are characterized by their extensive pre-training on large datasets, which enables impressive generalization and transfer learning capabilities. They are adaptable to downstream tasks through zero-shot or few-shot learning approaches, or can be fine-tuned to tackle specific problems. For instance, Subramanian et al. [27] explored the transfer learning capabilities of the Fourier Neural Operator (FNO) [19] by solving three classical PDEs and investigating its applicability across various physics, scales, and data availability in downstream tasks. The Unified Physics Solver (UPS) [28] extends this approach by covering a broader range of 1D and 2D PDEs, employing a pre-trained large language model for operator learning. Additionally, McCabe et al. [29] introduced a method to embed PDEs with varying physical properties into a shared embedding space, facilitating the simultaneous addressing

of multiple heterogeneous PDEs. Rahman et al. [30] proposed an attention mechanism tailored to the codomain of PDEs, enhancing the model’s ability to handle PDEs with varying dimensions.

Another burgeoning area within scientific machine learning focuses on leveraging in-context learning [31, 32, 33]. In this approach, models are prompted with multiple example pairs and are trained to make predictions on new query data based on patterns recognized during the training demonstrations. A notable implementation of this is the In-context Operator Network (ICON), which Yang et al. have explored in several studies [34, 35, 36]. ICON demonstrates operator learning by using example pairs that vary in parameters of the PDE and their corresponding solutions, thus enabling the network to predict solutions for new query data points.

3 Methodology

FMInt is pre-trained on a vast dataset comprising diverse solution trajectories of dynamical systems alongside their corresponding error correction terms. This extensive pre-training enables the model to effectively apply transfer learning, facilitating both zero-shot and few-shot inferences across a range of downstream applications. Our training methodology leverages in-context learning: the model observes examples of input data pairs and uses this context to make predictions on new, query data.

In solving Equation (1), one typically selects a numerical integration scheme that utilizes a large time step size, denoted as $k\Delta t$. For illustrative purposes, we consider the Euler method, which yields the following numerical simulation scheme:

$$\hat{\mathbf{u}}(t + k\Delta t) = \hat{\mathbf{u}}(t) + \mathbf{f}[\hat{\mathbf{u}}(t)] \cdot k\Delta t, \quad (3)$$

However, solving the dynamical system (1) with numerical scheme (3) and large step size unavoidably causes huge simulation errors, as we can see from the Taylor expansion

$$\mathbf{u}(t + k\Delta t) = \underbrace{\mathbf{u}(t) + \mathbf{f}[\mathbf{u}(t)] \cdot k\Delta t}_{\text{For Euler method}} + \underbrace{\sum_{n=2}^{\infty} \frac{1}{n!} \frac{d^n}{dt^n} \mathbf{u}(t) \cdot [k\Delta t]^n}_{\text{Error term err}_n}. \quad (4)$$

From (4), the error term $\sum_{n=2}^{\infty} \text{err}_n(k, \Delta t, \mathbf{u}(t))$ is non-negligible and limits the fast simulation of real-world dynamical systems. We call (3) solutions obtained by vanilla numerical integration schemes as “coarse solutions”. With coarse solutions as an initialization, our model is able to return highly accurate solution with fast inference time on a diverse set of dynamical systems, i.e.,

$$\hat{\mathbf{u}}_{k(n+1)} = \hat{\mathbf{u}}_{kn} + S(\mathbf{f}, \hat{\mathbf{u}}_{kn}, k\Delta t) + \text{FMInt}(\hat{\mathbf{u}}_{kn}; \Theta), \quad \hat{\mathbf{u}}_0 = \mathbf{c}_0, \quad n = 0, 1, \dots, \quad (5)$$

where Θ represents all the parameters of the foundation model.

We consider 500K ODEs that are commonly observed in important applications in engineering and science. To prepare our training data for the in-context learning of the model, we simulate each ODE system with coarse step size $k\Delta t$, and its corresponding error correction $\text{err}_{\hat{\mathbf{u}}_{kn}} = \mathbf{u}_{k(n+1)} - \hat{\mathbf{u}}_{kn} - S(\mathbf{f}, \hat{\mathbf{u}}_{kn}, k\Delta t)$. Therefore, the model takes the collection of “demos”, a query data point \mathbf{u}_t and outputs an error correction term err_t

$$\{\{\hat{\mathbf{u}}_0, \text{err}_0\}, \{\hat{\mathbf{u}}_k, \text{err}_k\}, \dots, \{\hat{\mathbf{u}}_{nk}, \text{err}_{nk}\}, \hat{\mathbf{u}}_t\} \rightarrow \text{err}_t. \quad (6)$$

Figure 1 illustrates our in-context learning approach for error correction, thus enabling fast inference of high accuracy solution from initial solution by classical numerical methods.

3.1 Data preparation

Data setup To pre-train our foundation model, we initially generate time series data from key dynamical systems that are prevalent across various applications. For each specified ODE, we create 1,000 variations with different parameters. Additionally, for each variation, we produce trajectories from 100 unique initial conditions. Consequently, for each ODE type, our dataset comprises the trajectories of 100,000 ODEs, each differentiated by varying coefficients and initial conditions.

1. **Decay equation.** We consider a first-order linear ODE of exponential decay processes

$$\frac{dN}{dt} = -\lambda N,$$

where N is the quantity that changes over time, and $\lambda \sim \text{Uniform}(0.2, 0.5)$ is the rate of decay.




Figure 1: In-context learning for error correction prediction. The model is prompted with N pairs of data, each consisting of a coarse numerical solution and its corresponding error correction, and a query coarse solution. Then, the model predicts the error in a forward pass without parameter update.

2. **Newton's Law of Cooling.** It describes how the temperature T of an object changes towards the environment temperature $T_{\text{env}} \sim \text{Uniform}(15, 25)$ over time, and $k \sim \text{Uniform}(0.1, 0.5)$ is the cooling rate coefficient.

$$\frac{dT}{dt} = -k(T - T_{\text{env}}),$$

3. **Lotka-Volterra system.** It is a pair of first-order, nonlinear differential equations used to describe the dynamics of biological systems

$$\begin{cases} \frac{dx}{dt} = \alpha x - \beta xy, \\ \frac{dy}{dt} = \delta xy - \gamma y, \end{cases}$$

where x is the number of prey, y is the number of predator, $\alpha \sim \text{Uniform}(0.1, 0.2)$ is the natural growing rate of prey in the absence of predators, $\beta \sim \text{Uniform}(0.04, 0.06)$ is the natural dying rate of prey due to predation, $\gamma \sim \text{Uniform}(0.6, 0.7)$ is the natural dying rate of predators in the absence of prey, and $\delta \sim \text{Uniform}(0.03, 0.04)$ is the rate at which predators increase by consuming prey.

4. **Van der Pol oscillator.** We consider a non-conservative oscillator with nonlinear damping, and it exhibits stable oscillatory solutions. The nonlinear damping also provides a rich framework for exploring complex behaviors such as bifurcations, chaos and pattern formation. The ODE reads as

$$\frac{d^2x}{dt^2} - \mu(1 - x^2)\frac{dx}{dt} + x = 0,$$

where x represents the state of system, $\mu \sim \text{Uniform}(0.5, 1.5)$ is a parameter that controls the strength of damping.

5. **Damped harmonic oscillator.** This equation is fundamental in physics and engineering for describing systems that exhibit oscillatory behavior with damping

$$\frac{d^2x}{dt^2} + 2\zeta\omega\frac{dx}{dt} + \omega^2x = 0,$$

where x represents the state of system, $\zeta \sim \text{Uniform}(0.01, 0.02)$ is the damping ratio, and $\omega \sim \text{Uniform}(5, 10)$ is the natural frequency of the undamped system.

The fine time step size Δt , the number of time steps k , the range of initial conditions (IC), and the numerical integration scheme are summarized in Table 1.

Table 1: Parameter setup

Name	k	Δt	IC (1d)	IC (2d)	Integration scheme
Exponential decay	10	0.05	(100, 200)	N/A	Euler
Newton’s law of cooling	10	0.05	(0, 80)	N/A	Euler
Lotka Volterra	200	0.005	(10, 20)	(2, 10)	RK4
Van der Pol	100	0.005	(−2.0, 2.0)	(−0.5, 0.5)	RK4
Damped	100	0.001	(−2.0, 2.0)	(−0.1, 0.1)	RK4

3.2 Model architecture

Inspired by the recent successes of large language models (LLMs) [24, 37, 38], we designed our model using a decoder-only transformer backbone [39]. This structure allows our model to predict the next “error” in a sequence by learning from previous examples, effectively correcting errors in each step of the simulation.

As illustrated in Table 2, we structure our input using three types of tokens: condition tokens, Quantity of Interest (QoI) tokens, and query tokens, following [34]. For the i -th example, each token comprises a key-value pair. The key represents the coarse time step, denoted as $ik\Delta t$. The values for the condition and QoI tokens correspond to the coarse solution and error correction term at time $ik\Delta t$, respectively, while the value for query tokens is populated with zero.

Each token undergoes transformation through a shared embedding layer to obtain a representative embedding vector. This vector is then concatenated with a learnable positional embedding to maintain temporal context. To preserve order invariance among key-value pairs, all tokens of the same type within an example share the same positional encoding. The concatenated vectors are fed into the transformer model, and the transformer’s output is subsequently directed to a prediction head for final error prediction. Our model is trained using the mean squared error as the loss function.

A major challenge in training decoder-only transformer models is the implementation of masking. Specifically, when predicting the QoI, the model must consider the conditions and QoI values from preceding examples. This requirement stems from the fact that all demonstration examples pertain to the identical ODE, differing only in their initial conditions. Furthermore, predictions of QoI are independent and remain unaffected by the order of the tokens. To effectively manage these constraints, we employ a specialized transformer masking technique, as described in Yang et al. [36].

Table 2: The tokens for a example.

key	condition				QoI				query			
	0	$k\Delta t$...	$(n-1)k\Delta t$	0	$k\Delta t$...	$(n-1)k\Delta t$	0	$k\Delta t$...	$(n-1)k\Delta t$
value	$\hat{u}(0)$	$\hat{u}(t_1)$...	$\hat{u}(t_{n-1})$	$\text{err}_{\hat{u}}(0)$	$\text{err}_{\hat{u}}(t_1)$...	$\text{err}_{\hat{u}}(t_{n-1})$	0	0	...	0
	$\hat{v}(0)$	$\hat{v}(t_1)$...	$\hat{v}(t_{n-1})$	$\text{err}_{\hat{v}}(0)$	$\text{err}_{\hat{v}}(t_1)$...	$\text{err}_{\hat{v}}(t_{n-1})$	0	0	...	0

4 Numerical Experiments

In this section, we aim to show the effectiveness of FMint in obtaining high accuracy solution based on initial coarse solutions.

4.1 Model set-up

As a decoder-only transformer model, FMint is configured with approximately 15.8 million parameters. The model features six heads for multi-head attention, with an input/output dimension of 256 for each layer. Additionally, the dimension for the query, key, and value of each token is set to 256, while the hidden dimension of the feedforward networks is 1024. All experiments were conducted on a NVIDIA A100 GPU with 80 GB of memory. We use AdamW optimizer with a warmup-cosine-decay schedule, with peak learning rate $1\text{e-}4$ and 300,000 training steps. The Adam β_1 and Adam β_2 are 0.9 and 0.999, respectively. The Adam weight decay is set to $1\text{e-}4$.

Table 3: Relative errors and RMSE errors

Name	corrected coarse-grained ODE		coarse-grained ODE	
	Avg. rel. err.	RMSE	Avg. rel. err.	RMSE
Exponential decay	3.49e-2	7.80e-3	3.64e-1	6.47
Newton’s law of cooling	1.67e-4	5.57e-3	2.37e-2	1.31
Lotka Volterra	1.61e-3	5.85e-3	3.71e-1	1.54
damped harmonic oscillator	4.54e-2	1.22e-2	1.93	2.39
Van der Pol oscillator	2.02e-2	1.17e-2	6.83e-1	1.56

4.2 In-distribution ODE

For all test examples, we compute the mean relative errors (MRE) and average L_2 errors for coarse-grained ODE compared to fine-grained ODE solutions with five demonstrations:

$$\text{MRE} = \frac{1}{KN} \sum_{i=1}^N \sum_k^K \frac{|\hat{u}_k + \hat{\text{err}}_k - u_k|}{|u_k|} \quad (7)$$

and

$$\text{RMSE} = \sqrt{\frac{1}{KN} \sum_{i=1}^N \sum_k^K \|\hat{u}_k + \hat{\text{err}}_k - u_k\|_2^2} \quad (8)$$

where $\hat{\text{err}}$ is the predicted error. As shown in Table 3, the MRE and RMSE of the corrected coarse-grained ODE solutions are significantly lower than those of the coarse-grained ODE solutions. This demonstrates the effectiveness of FMint in enhancing the accuracy of coarse-grained ODE solutions. Additional results are shown in Figure 2 and Figure 3.

5 Conclusion

In conclusion, we presented FMint, a novel generative pre-trained model that successfully integrates the strengths of human-designed algorithms and data-driven methods for the simulation of dynamical systems. By leveraging a vast dataset of dynamical systems with varying parameters and conditions, FMint demonstrates a significant advancement in the ability to generate accurate simulations swiftly and efficiently.

Our model’s incorporation of in-context learning, based on a decoder-only transformer architecture, allows for exceptional adaptability and generalization across a diverse range of dynamical systems. This approach not only surpasses the performance of traditional numerical integration methods, which often struggle with flexibility and domain-specific constraints, but also provides a robust solution to the challenges faced by other neural network-based methods, such as NeurVec, which are limited by costly retraining and generalization capabilities.

Furthermore, FMint’s proficiency in zero-shot and few-shot learning makes it as a foundational tool in scientific computing, capable of solving problems with unprecedented accuracy and speed. This makes it a powerful tool for real-world applications where fast and reliable simulation of dynamical systems is critical.


As we continue to develop and refine FMint, future research will focus on expanding its applicability to even more complex systems and exploring the potential synergies with other machine learning techniques. The ongoing evolution of foundation models like FMint will play a crucial role in the advancement of scientific machine learning, offering possibilities for broader applications.

References


- [1] Roger Temam. *Infinite-dimensional dynamical systems in mechanics and physics*, volume 68. Springer Science & Business Media, 2012.
- [2] James D Meiss. *Differential dynamical systems*. SIAM, 2007.
- [3] Denis L Blackmore, Valeriy Hr Samoylenko, et al. *Nonlinear dynamical systems of mathematical physics: spectral and symplectic integrability analysis*. World Scientific, 2011.

- [4] Tamás Tél, Alessandro de Moura, Celso Grebogi, and György Károlyi. Chemical and biological activity in open flows: A dynamical system approach. *Physics reports*, 413(2-3):91–196, 2005.
- [5] Christian Vidal and Adolphe Pacault. *Non-Equilibrium Dynamics in Chemical Systems: Proceedings of the International Symposium, Bordeaux, France, September 3–7, 1984*, volume 27. Springer Science & Business Media, 2012.
- [6] Vasile Marinca and Nicolae Herisanu. *Nonlinear dynamical systems in engineering: Some approximate approaches*. Springer Science & Business Media, 2012.
- [7] Stephen Wiggins. The dynamical systems approach to lagrangian transport in oceanic flows. *Annu. Rev. Fluid Mech.*, 37:295–328, 2005.
- [8] Rafal Goebel, Ricardo G Sanfelice, and Andrew R Teel. Hybrid dynamical systems. *IEEE control systems magazine*, 29(2):28–93, 2009.
- [9] Dominique Guegan. Chaos in economics and finance. *Annual Reviews in Control*, 33(1):89–93, 2009.
- [10] J Dong, D Zhang, and A Nagurney. A projected dynamical systems model of general financial equilibrium with stability analysis. *Mathematical and computer Modelling*, 24(2):35–44, 1996.
- [11] 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.
- [12] Joan Bruna, Benjamin Peherstorfer, and Eric Vanden-Eijnden. Neural galerkin schemes with active learning for high-dimensional evolution equations. *Journal of Computational Physics*, 496:112588, 2024.
- [13] George Em Karniadakis, Ioannis G Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.
- [14] Tianqi Cui, Tom Bertalan, Nelson Ndahiro, Pratik Khare, Michael Betenbaugh, Costas Maranas, and Ioannis G Kevrekidis. Data-driven and physics informed modeling of chinese hamster ovary cell bioreactors. *Computers & Chemical Engineering*, 183:108594, 2024.
- [15] Mario De Florio, Ioannis G Kevrekidis, and George Em Karniadakis. Ai-lorenz: A physics-data-driven framework for black-box and gray-box identification of chaotic systems with symbolic regression. *arXiv preprint arXiv:2312.14237*, 2023.
- [16] Senwei Liang and Haizhao Yang. Finite expression method for solving high-dimensional partial differential equations. *arXiv preprint arXiv:2206.10121*, 2022.
- [17] Zezheng Song, Maria K Cameron, and Haizhao Yang. A finite expression method for solving high-dimensional committor problems. *arXiv preprint arXiv:2306.12268*, 2023.
- [18] Zezheng Song, Chunmei Wang, and Haizhao Yang. Finite expression method for learning dynamics on complex networks. *arXiv preprint arXiv:2401.03092*, 2024.
- [19] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations (2020). *arXiv preprint arXiv:2010.08895*, 2010.
- [20] Yong Zheng Ong, Zuowei Shen, and Haizhao Yang. Iae-net: Integral autoencoders for discretization-invariant learning. *arXiv preprint arXiv:2203.05142*, 2022.
- [21] Shuhao Cao. Choose a transformer: Fourier or galerkin. *Advances in neural information processing systems*, 34:24924–24940, 2021.
- [22] Zijie Li, Kazem Meidani, and Amir Barati Farimani. Transformer for partial differential equations’ operator learning. *arXiv preprint arXiv:2205.13671*, 2022.
- [23] Zhongzhan Huang, Senwei Liang, Hong Zhang, Haizhao Yang, and Liang Lin. On fast simulation of dynamical system with neural vector enhanced numerical solver. *Scientific Reports*, 13(1):15254, 2023.
- [24] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [25] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International conference on machine learning*, pages 8821–8831. Pmlr, 2021.


- [26] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [27] 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. *Advances in Neural Information Processing Systems*, 36, 2024.
- [28] Junhong Shen, Tanya Marwah, and Ameet Talwalkar. Ups: Towards foundation models for pde solving via cross-modal adaptation. *arXiv preprint arXiv:2403.07187*, 2024.
- [29] Michael McCabe, Bruno Régaldou-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.
- [30] Md Ashiqur Rahman, Robert Joseph George, Mogab Elleithy, Daniel Leibovici, Zongyi Li, Boris Bonev, Colin White, Julius Berner, Raymond A Yeh, Jean Kossaifi, et al. Pretraining codomain attention neural operators for solving multiphysics pdes. *arXiv preprint arXiv:2403.12553*, 2024.
- [31] Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. A survey on in-context learning. *arXiv preprint arXiv:2301.00234*, 2022.
- [32] Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. An explanation of in-context learning as implicit bayesian inference. *arXiv preprint arXiv:2111.02080*, 2021.
- [33] Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, et al. In-context learning and induction heads. *arXiv preprint arXiv:2209.11895*, 2022.
- [34] Liu Yang, Siting Liu, Tingwei Meng, and Stanley J Osher. In-context operator learning for differential equation problems. *arXiv preprint arXiv:2304.07993*, 2023.
- [35] Liu Yang and Stanley J Osher. Pde generalization of in-context operator networks: A study on 1d scalar nonlinear conservation laws. *arXiv preprint arXiv:2401.07364*, 2024.
- [36] Liu Yang, Tingwei Meng, Siting Liu, and Stanley J Osher. Prompting in-context operator learning with sensor data, equations, and natural language. *arXiv preprint arXiv:2308.05061*, 2023.
- [37] Arun James Thirunavukarasu, Darren Shu Jeng Ting, Kabilan Elangovan, Laura Gutierrez, Ting Fang Tan, and Daniel Shu Wei Ting. Large language models in medicine. *Nature medicine*, 29(8):1930–1940, 2023.
- [38] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.



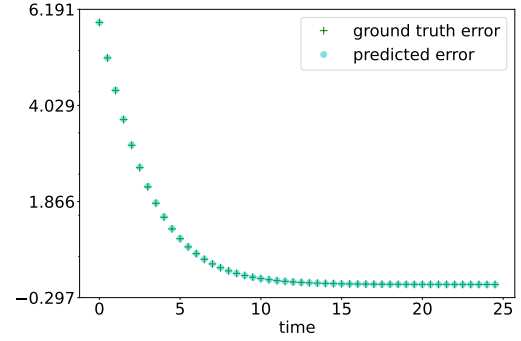
(a) ODE of exponential decay




(b) Error prediction of exponential decay




(c) ODE of Newton's law of cooling



(d) Error prediction of Newton's law of cooling




(e) ODE of Lotka Volterra




(f) Error prediction of Lotka Volterra


Figure 2: The predicted fine-grained ODE solutions and ground truth fine-grained ODE solutions for exponential decay (1D), Newton's law of cooling (1D) and Lotka Volterra equation (2D) are shown in (a), (c) and (e). The predicted errors and ground truth error are shown in (b), (d) and (f).




(a) ODE of Van der Pol oscillators



(b) Error prediction of Van der Pol oscillators



(c) ODE of damped harmonic oscillators



(d) Error prediction of damped harmonic oscillators

Figure 3: The predicted fine-grained ODE solutions and ground truth fine-grained ODE solutions for Van der Pol oscillators (2D) and damped harmonic oscillators (2D) are shown in (a) and (c). The predicted errors and ground truth error are shown in (b) and (d).