
Beyond Closure Models: Learning Chaotic-Systems via Physics-Informed Neural Operators

Chuwei Wang¹, Julius Berner¹, Zongyi Li¹, Di Zhou², Jiayun Wang¹,
Jane Bae², Anima Anandkumar^{1*}

¹ Department of Computing and Mathematical Sciences, Caltech

² Graduate Aerospace Laboratories, Caltech

{chuweiw, jberner, zongyili, dizhou, peterw, jbae, anima}@caltech.edu

Abstract

Accurately predicting the long-term behavior of chaotic systems is crucial for various applications such as climate modeling. However, achieving such predictions typically requires iterative computations over a dense spatiotemporal grid to account for the unstable nature of chaotic systems, which is expensive and impractical in many real-world situations. An alternative approach to such a fully-resolved simulation (FRS) is using a coarse grid and then correcting its errors through a *closure model*, which approximates the overall information from fine scales not captured in the coarse-grid simulation. Recently, ML approaches have been used for closure modeling, but they typically require a large number of expensive FRS training samples. In this work, we prove an even more fundamental limitation, i.e., the standard approach to learning closure models suffers from a large approximation error no matter how expressive the model is, and it stems from the non-uniqueness of the mapping. We also prove that other existing methods leveraging history information and randomness can neither resolve this limitation. We propose an alternative end-to-end learning approach using a physics-informed neural operator (PINO) that overcomes this limitation by not using a closure model or a coarse-grid solver. We first train the PINO model on data from a coarse-grid solver and then fine-tune it with (a small amount of) FRS and physics-based losses on a fine grid. The discretization-free nature of neural operators means that they do not suffer from the restriction of a coarse grid that closure models face, and they can provably approximate the long-term statistics of chaotic systems. In our experiments, our PINO model achieves a 120x speedup compared to FRS with a relative error $\sim 5\%$. In contrast, the closure model coupled with a coarse-grid solver is 58x slower than PINO while having a higher error $\sim 205\%$ when trained on the same FRS dataset.

1 Introduction

Predicting long-term behavior is an important task in many physical systems, e.g., climate modeling, aircraft design, and plasma evolution in nuclear fusion [1, 2, 3, 4, 5]. This can be framed as estimating statistics of a system in its dynamical equilibrium through numerical simulations. One major challenge is that many physical systems are chaotic and thus have extreme sensitivity to perturbations [6, 7, 8, 9]. To account for the unstable nature of chaotic systems, high-fidelity simulations have to be carried out on extremely fine spatiotemporal grids to make discretization errors sufficiently small so that the overall error along the trajectory does not grow rapidly. This makes fully-resolved simulations (FRS), e.g. direct numerical simulations (DNS) in turbulence, prohibitively expensive in terms of both computation time and memory [10, 11].

*Correspondence to : Anima Anandkumar <anima@caltech.edu>.

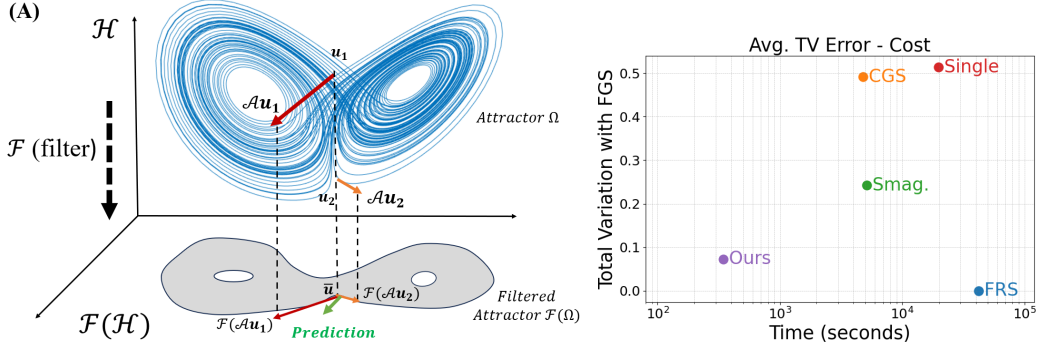


Figure 1: **Left:** Many points (e.g., u_1 and u_2) of the ground-truth attractor Ω (i.e., equilibrium state, blue) map to the same filtered (e.g., downsampled) value \bar{u} (gray), making it impossible for the closure model to identify the correct dynamics ($\mathcal{F}(\mathcal{A}u_1)$ and $\mathcal{F}(\mathcal{A}u_2)$) in the filtered space $\mathcal{F}(\mathcal{H})$. By minimizing the loss function, the model learns to predict an average of these multiple choices (green arrow), which leads the simulation to wrongly diverge from the filtered attractor. **Right:** Total variation distance from ground-truth invariant measure versus computation cost. ‘FRS’ (blue line): gold-standard fully-resolved simulations. ‘CGS’: coarse-grid simulation without closure model. ‘Smag.’: Smagorinsky model. ‘Single’: learning-based single-state model. Our method is the fastest and closest to ground truth (‘FRS’) among all coarse-grid methods (‘CGS’, ‘Smag.’, ‘Single’).

Given the computation cost of FRS in chaotic dynamics and the fact that the ultimate goal is to evaluate the long-term statistics instead of tracking any individual trajectory, many works have been exploring ways to give a good estimate of such statistics with simulations only conducted on coarse spatial grids, e.g., large-eddy simulation (LES) [12, 13] for turbulent flows. To correct the errors introduced by coarse-grid solvers, a popular approach is to use a *closure model*, an additional term in the original dynamics, in conjunction with the solver [14, 15, 16]. This approach is also known as coarse-grained modeling [17] or renormalization groups in some disciplines [18].

The traditional framework for designing closure models is based on physical intuition, which requires substantial domain expertise or derived by mathematical simplification under strong modeling assumptions. Hence, such approaches are typically inaccurate for real-world systems [19, 20, 21]. This past decade has also witnessed extensive development of machine-learning methods for closure modeling [22, 23, 24, 25, 26]. However, these learning-based closure models [27, 28, 29, 30] often rely on a large amount of high-fidelity training data generated from expensive FRS, which may even be impossible to generate for many problems of interest.

Our Approach: In this work, we provide a new theoretical understanding of estimating long-term statistics of a chaotic system. We formally prove in Theorem 2.1 that for generic problems, previous learning methods based on closure models are fundamentally ill-posed since they are constrained to be on the same coarse grid as the solver, and they cannot accurately approximate the underlying chaotic system. Specifically, we prove that the mapping that closure models attempt to learn in a reduced space (coarse grid) is non-unique, i.e., there are multiple potential outputs for a given input (fig. 1, left). Hence, the standard approach to learning closure models under such non-uniqueness results in the average of possible outputs, and that cannot accurately approximate the long-term statistics of a chaotic system, no matter how expressive the closure model is. We also prove that other attempts leveraging history information or randomness can neither resolve this limitation.

We further propose an alternative end-to-end ML framework to mitigate this non-unique issue of the previous scheme, i.e. an additional closure model in coarse-grid simulations. We remove the constraint that the learned model is on a coarse grid and instead employ a grid-free approach to learning. It is based on neural operators [31, 32], which learn mappings between function spaces. In a neural operator, inputs of different resolutions are viewed as different discretizations of the same function on a continuous domain, thus ensuring consistency between coarse and fine grids. The neural operator, targeting the solution operator directly, is first pre-trained with data from cheap coarse-grid solvers and a small amount of FRS data, then trained by minimizing physics-based loss defined on a fine grid. The physics-based losses on a fine grid reduce the FRS data requirement and improve generalization, in line with what has been seen in prior works on physics-informed learning [33].

We prove in Theorem A.1 that a neural operator that approximates the underlying ground-truth solution operator can provide sufficiently accurate estimates of the long-term statistics of a chaotic

system, and there is no catastrophic build-up of errors over long rollouts. We test the performance of our approach in several instances from fluid dynamics. Our method achieves a 120x speedup with a relative error $\sim 5\%$ compared to FRS. In contrast, the closure model coupled with a coarse-grid solver is 58x slower than ours while having a much higher error $\sim 205\%$. More details about the speed-accuracy performance of different approaches are shown in figure 1. The closure model has a significantly higher error under our setup compared to prior works [30, 29, 34] because we assume that only a few FRS samples are available for training, both for PINO and the closure model. In fact, we show that to mitigate the non-unique issue of closure models mentioned above, previous methods have to resort to the closeness between the limit distribution in the original dynamics and the empirical measure of training data. An illustrative comparison between our method and representative existing methods can be found in Table 2. Our contributions are summarized in Appendix A.4

2 Background and Existing Methods

Problem Background Consider an evolution partial differential equation (PDE) that governs a (nonlinear) dynamical system in the function space,

$$\begin{cases} \partial_t u(x, t) = \mathcal{A}u(x, t) \\ u(x, 0) = u_0(x), u_0 \in \mathcal{H}, \end{cases} \quad (1)$$

where u_0 is the initial value and \mathcal{H} is a function space containing functions of interests. We denote by $S(t)$ the corresponding semi-group. We refer to $\{S(t)u_0\}_{t \geq 0}$ as *trajectory from u_0* . In many physics systems, there exists a global *attractor* defined as a subset of \mathcal{H} towards which all trajectories converge over time [35, 36]. The *invariant measure* μ^* , supported on the attractor, is the time average over any trajectory, independent of initial value as long as the system is ergodic. Intuitively, μ^* captures the system’s long-term behavior when it reaches a dynamical equilibrium. The *long-term statistics* are expectations of functionals on the invariant measure. See Appendix A.2 for more details.

Coarse-grid Simulation and Closure Modeling Many works have been exploring ways to *give good estimations of statistics with simulations only conducted on coarse grids*. We will refer to this approach as coarse-grid simulation (CGS). CGS could be viewed as evolving a filtered function \bar{u} defined as $\bar{u} = \mathcal{F}u$, where \mathcal{F} is a linear filtering operator. For instance, \mathcal{F} is a spatial convolution for cases like down-sampling in the finite difference method and Fourier-mode truncation in the spectral method. Theoretically, the evolution of \bar{u} is governed by $\partial_t \bar{u} = \mathcal{F}\mathcal{A}u = \mathcal{A}\bar{u} + (\mathcal{F}\mathcal{A} - \mathcal{A}\mathcal{F})u$. However, the commutator $(\mathcal{F}\mathcal{A} - \mathcal{A}\mathcal{F})u$ is intractable if restricted to coarse grids since u is underresolved. To account for the effect of small scales not captured by coarse grids, in many CGS methods an adjusting term $clos(\bar{u}; \theta)$ (θ denotes the model parameters), known as *closure model*, is added to the equation as a tractable surrogate model of $(\mathcal{F}\mathcal{A} - \mathcal{A}\mathcal{F})u$. The CGS trajectory is

$$\begin{cases} \partial_t v(x, t) = \mathcal{A}v(x, t) + clos(v; \theta), x \in D' \\ v(x, 0) = \bar{u}_0(x), \bar{u}_0 \in \mathcal{F}(\mathcal{H}), \end{cases} \quad (2)$$

and the statistics are estimated as the time average of the corresponding functionals with $v(\cdot, t)$ input.

Learning-based closure model: In recent years, there has been a growing interest in leveraging machine learning tools to design closure models (see [24] for survey). The most popular approach is to learn a single-state closure model $clos(v; \theta)$ with various ansatz and neural network architecture, trained by minimizing an a priori loss function aiming at fitting the commutator [24],

$$J_{ap}(\theta; \mathcal{D}) = \frac{1}{|\mathcal{D}|} \sum_{i \in \mathcal{D}} \|clos(\bar{u}_i; \theta) - (\mathcal{F}\mathcal{A} - \mathcal{A}\mathcal{F})u_i\|^2, \quad (3)$$

where the training data u_i come from snapshots of FRS trajectories, i.e., $S(t)u_0$ for particular t . This methodology is problematic when considering the input and output. Due to the dimension reduction of filter \mathcal{F} , multiple FRS snapshots $u_i \in \mathcal{H}$ map to the same state in the reduced space $\mathcal{F}(\mathcal{H})$. However, the filtered vector field $\mathcal{A}u$ differs across these u_i . The closure model must assign a unique direction in $\mathcal{F}(\mathcal{H})$ for each state \bar{u} . Minimizing the training loss (3) often results in predicting an averaged vector field, which does not make much sense. See Fig. (1). Unlike inverse problems or linear regression, predictions here must follow the manifold defined by physical trajectories.

There are also other variants that leverages history information [34, 37] (a closure model whose input is $\{\bar{u}(\cdot, t-s)\}_{0 < s \leq t_0}$ at the moment t , where t_0 is a model parameter) or propose to employ a stochastic closure model [38, 39]. However, we have the following results.

Theorem 2.1. (i) In general, the mapping of closure models $\bar{u} \rightarrow (\mathcal{F}\mathcal{A} - \mathcal{A}\mathcal{F})u$ is not well-defined. Consequently, the approximation error has a lower bound independent of the model complexity.
(ii) For any u and finite τ , there exist **infinite** $u' \in \mathcal{H}$ such that $\mathcal{F}S(t)u' = \mathcal{F}S(t)u$ for all $t \in [0, \tau)$.
(iii) One cannot obtain the best approximation of μ^* among distributions supported in the reduced space if there is randomness in the evolution of dynamics.

The proof is in Appendix D. From the second claim, even with history-aware models, under-determinacy remains in chaotic systems because the reduced-space history does not fully determine future states. For the third claim, parameters associated with randomness in stochastic models will tend to zero after ideal training, making randomness unnecessary.

Given the results above, it might seem surprising since these methods still achieve competitive performance. We show that their empirical result depends heavily on having abundant FRS training data (see Appendix A.5 for details). If enough FRS data were available, we could compute the statistics directly, negating the need for a closure model or coarse-grid simulations.

3 Methodology: Physics-Informed Operator Learning

From previous sections, we see that restricting the learning object in the filtered space and explicitly learning the closure model would always suffer from the non-uniqueness of this target. In light of that, we propose to extend the learning task into function space \mathcal{H} and directly deal with the solution operator $S(t)$ of PDE governing the dynamics, which is a well-defined mapping. We adopt Fourier Neural Operator (FNO) [31] to learn the mapping $u \rightarrow \{S(t)u\}_{t \in [0, h]}$, where h is a model parameter.

To overcome the lack of FRS training data in realistic situations, we adopt physics-informed methodologies[40] to remove the reliance on data. To be specific, the operator model \mathcal{G}_θ is trained by minimizing the physics-informed loss function

$$J_{pde}(\theta; \mathfrak{D}) = \frac{1}{|\mathfrak{D}|} \sum_{i \in \mathfrak{D}} \|(\partial_t - \mathcal{A})\mathcal{G}_\theta u_{0i}(x)\|_{L^2(\Omega \times [0, h])}, \quad (4)$$

where the initial values u_{0i} in the loss function could be any fine-grid functions and do not have to come from FRS trajectories. Ω is the spatial domain of these functions. In practice, the optimization of physics-informed loss is hard [41, 42]. To face these challenges, we pre-train the model via supervised learning with a data loss function to achieve a good initialization of the model parameters for J_{pde} optimization. A detailed introduction of neural operator is in Appendix A.6, the formalized algorithm and its implementation details can be found in Appendix G. A theoretical guarantee for estimating statistics is presented in Appendix A.7.

Experiment Results To verify our method, we compare it with gold-standard fully-resolved simulations, coarse-grid simulations with numerical closure models [14] and ML-based single-state closure models which are trained with same amount of FRS data as our model, namely 110 snapshots from a single trajectory. We test on 1D Kuramoto–Sivashinsky equation and 2D Navier–Stokes equation. We check various statistics including energy spectrum, auto-correlation, variance, velocity and vorticity density, kinetic energy and dissipation rate. See Table 1. More details about the experiment setup, the statistics considered, and full results for all equations and statistics are in Appendices F to H.

Table 1: **Experiment Results for Navier-Stokes Equation. Left:** Errors on different statistics, i.e., average total variation (‘Avg. TV’), energy spectrum (‘Energy’), TV error for vorticity distribution (‘Vorticity’), and velocity variance (‘Variance’). Percentages refer to average relative errors. Other numbers refer to TV distances (ranging $[0, 1]$) between ground truth and prediction. **Right:** Comparison of the inference time (seconds) of one trajectory for $t \in [0, 100]$. Best results are marked **bold**.

Method	Avg. TV	Energy	Vorticity	Variance	FRS	
CGS (No closure)	0.4914	178.4651%	0.1512	253.4234%	CGS (No closure)	39.70
Smagorinsky [14]	0.2423	52.9511%	0.0483	20.1740%	Smagorinsky	4.81
Single-state [30]	0.5137	205.3709%	0.1648	298.2027%	Single-state	18.57
Our Method	0.0726	5.3276%	0.0091	2.8666%	Ours	0.32

References

- [1] JAS Lima, R Silva Jr, and Janilo Santos. Plasma oscillations and nonextensive statistics. *Physical Review E*, 61(3):3260, 2000.
- [2] Gregory Flato, Jochem Marotzke, Babatunde Abiodun, Pascale Braconnot, Sin Chan Chou, William Collins, Peter Cox, Fatima Driouech, Seita Emori, Veronika Eyring, et al. Evaluation of climate models. In *Climate change 2013: the physical science basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, pages 741–866. Cambridge University Press, 2014.
- [3] Stephen H Schneider and Robert E Dickinson. Climate modeling. *Reviews of Geophysics*, 12(3):447–493, 1974.
- [4] Jeffrey P Slotnick, Abdollah Khodadoust, Juan Alonso, David Darmofal, William Gropp, Elizabeth Lurie, and Dimitri J Mavriplis. CFD vision 2030 study: a path to revolutionary computational aerosciences. Technical Report CR-2014–21817, NASA, 2014.
- [5] David M Wootton and David N Ku. Fluid mechanics of vascular systems, diseases, and thrombosis. *Annual Review of Biomedical Engineering*, 1(1):299–329, 1999.
- [6] Julio M Ottino. Mixing, chaotic advection, and turbulence. *Annual Review of Fluid Mechanics*, 22(1):207–254, 1990.
- [7] Gerald Jay Sussman and Jack Wisdom. Chaotic evolution of the solar system. *Science*, 257(5066):56–62, 1992.
- [8] Henri Korn and Philippe Faure. Is there chaos in the brain? ii. experimental evidence and related models. *Comptes Rendus Biologies*, 326(9):787–840, 2003.
- [9] Steven H Strogatz. *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering*. CRC press, 2018.
- [10] Kiran Ravikumar, David Appelhans, and PK Yeung. Gpu acceleration of extreme scale pseudo-spectral simulations of turbulence using asynchronism. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–22, 2019.
- [11] Tapio Schneider, João Teixeira, Christopher S Bretherton, Florent Brient, Kyle G Pressel, Christoph Schär, and A Pier Siebesma. Climate goals and computing the future of clouds. *Nature Climate Change*, 7(1):3–5, 2017.
- [12] Javier Jimenez and Robert D Moser. Large-eddy simulations: where are we and what can we expect? *AIAA journal*, 38(4):605–612, 2000.
- [13] Stephen B Pope. Ten questions concerning the large-eddy simulation of turbulent flows. *New Journal of Physics*, 6(1):35, 2004.
- [14] Joseph Smagorinsky. General circulation experiments with the primitive equations: I. the basic experiment. *Monthly Weather Review*, 91(3):99–164, 1963.
- [15] Yupeng Zhang and Kaushik Bhattacharya. Iterated learning and multiscale modeling of history-dependent architected metamaterials. *arXiv preprint arXiv:2402.12674*, 2024.
- [16] Burigede Liu, Eric Ocegueda, Margaret Trautner, Andrew M Stuart, and Kaushik Bhattacharya. Learning macroscopic internal variables and history dependence from microscopic models. *Journal of the Mechanics and Physics of Solids*, 178:105329, 2023.
- [17] Valentina Tozzini. Coarse-grained models for proteins. *Current Opinion in Structural Biology*, 15(2):144–150, 2005.
- [18] Kenneth G Wilson. Renormalization group and strong interactions. *Physical Review D*, 3(8):1818, 1971.

- [19] Charles Meneveau and Joseph Katz. Scale-invariance and turbulence models for large-eddy simulation. *Annual Review of Fluid Mechanics*, 32(1):1–32, 2000.
- [20] Robert D Moser, Sigfried W Haering, and Gopal R Yalla. Statistical properties of subgrid-scale turbulence models. *Annual Review of Fluid Mechanics*, 53:255–286, 2021.
- [21] Di Zhou and H Jane Bae. Sensitivity analysis of wall-modeled large-eddy simulation for separated turbulent flow. *Journal of Computational Physics*, 506:112948, 2024.
- [22] Karthik Duraisamy, Gianluca Iaccarino, and Heng Xiao. Turbulence modeling in the age of data. *Annual Review of Fluid Mechanics*, 51:357–377, 2019.
- [23] Karthik Duraisamy. Perspectives on machine learning-augmented reynolds-averaged and large eddy simulation models of turbulence. *Physical Review Fluids*, 6(5):050504, 2021.
- [24] Benjamin Sanderse, Panos Stinis, Romit Maulik, and Shady E Ahmed. Scientific machine learning for closure models in multiscale problems: a review. *arXiv preprint arXiv:2403.02913*, 2024.
- [25] Suryanarayana Maddu, Scott Weady, and Michael J Shelley. Learning fast, accurate, and stable closures of a kinetic theory of an active fluid. *Journal of Computational Physics*, 504:112869, 2024.
- [26] Varun Shankar, Vedant Puri, Ramesh Balakrishnan, Romit Maulik, and Venkatasubramanian Viswanathan. Differentiable physics-enabled closure modeling for burgers’ turbulence. *Machine Learning: Science and Technology*, 4(1):015017, 2023.
- [27] Masataka Gamahara and Yuji Hattori. Searching for turbulence models by artificial neural network. *Physical Review Fluids*, 2(5):054604, 2017.
- [28] Andrea Beck, David Flad, and Claus-Dieter Munz. Deep neural networks for data-driven LES closure models. *Journal of Computational Physics*, 398:108910, 2019.
- [29] Romit Maulik, Omer San, Adil Rasheed, and Prakash Vedula. Subgrid modelling for two-dimensional turbulence using neural networks. *Journal of Fluid Mechanics*, 858:122–144, 2019.
- [30] Yifei Guan, Ashesh Chattopadhyay, Adam Subel, and Pedram Hassanzadeh. Stable a posteriori LES of 2D turbulence using convolutional neural networks: Backscattering analysis and generalization to higher re via transfer learning. *Journal of Computational Physics*, 458:111090, 2022.
- [31] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.
- [32] 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.
- [33] 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. *ACM/JMS Journal of Data Science*, 2021.
- [34] Chao Ma, Jianchun Wang, et al. Model reduction with memory and the machine learning of dynamical systems. *arXiv preprint arXiv:1808.04258*, 2018.
- [35] Roger Temam. *Infinite-dimensional dynamical systems in mechanics and physics*, volume 68. Springer Science & Business Media, 2012.
- [36] John Milnor. On the concept of attractor. *Communications in Mathematical Physics*, 99:177–195, 1985.

- [37] Qian Wang, Nicolò Ripamonti, and Jan S Hesthaven. Recurrent neural network closure of parametric pod-galerkin reduced-order models based on the mori-zwanzig formalism. *Journal of Computational Physics*, 410:109402, 2020.
- [38] Anudhyan Boral, Zhong Yi Wan, Leonardo Zepeda-Núñez, James Lottes, Qing Wang, Yi-fan Chen, John Roberts Anderson, and Fei Sha. Neural ideal large eddy simulation: Modeling turbulence with neural stochastic differential equations. *arXiv preprint arXiv:2306.01174*, 2023.
- [39] Fei Lu, Kevin K Lin, and Alexandre J Chorin. Data-based stochastic model reduction for the kuramoto–sivashinsky equation. *Physica D: Nonlinear Phenomena*, 340:46–57, 2017.
- [40] 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.
- [41] Pratik Rathore, Weimu Lei, Zachary Frangella, Lu Lu, and Madeleine Udell. Challenges in training pinns: A loss landscape perspective. *arXiv preprint arXiv:2402.01868*, 2024.
- [42] Chuwei Wang, Shanda Li, Di He, and Liwei Wang. Is l^2 physics informed loss always suitable for training physics informed neural network? *Advances in Neural Information Processing Systems*, 35:8278–8290, 2022.
- [43] William C Reynolds. The potential and limitations of direct and large eddy simulations. In *Whither Turbulence? Turbulence at the Crossroads: Proceedings of a Workshop Held at Cornell University, Ithaca, NY, March 22–24, 1989*, pages 313–343. Springer, 2005.
- [44] Haecheon Choi and Parviz Moin. Grid-point requirements for large eddy simulation: Chapman’s estimates revisited. *Physics of Fluids*, 24(1), 2012.
- [45] Justin Sirignano and Jonathan F MacArt. Dynamic deep learning LES closures: Online optimization with embedded DNS. *arXiv preprint arXiv:2303.02338*, 2023.
- [46] Sergei P Kuznetsov. Dynamical chaos and uniformly hyperbolic attractors: from mathematics to physics. *Physics-USpekhi*, 54(2):119, 2011.
- [47] Simona Dinicola, Fabrizio D’Anselmi, Alessia Pasqualato, Sara Proietti, Elisabetta Lisi, Alessandra Cucina, and Mariano Bizzarri. A systems biology approach to cancer: fractals, attractors, and nonlinear dynamics. *Omics: a journal of integrative biology*, 15(3):93–104, 2011.
- [48] Sui Huang, Ingemar Ernberg, and Stuart Kauffman. Cancer attractors: a systems view of tumors from a gene network dynamics and developmental perspective. In *Seminars in cell & developmental biology*, volume 20, pages 869–876. Elsevier, 2009.
- [49] Alfredo Medio and Marji Lines. *Nonlinear dynamics: A primer*. Cambridge University Press, 2001.
- [50] Jacob A Langford and Robert D Moser. Optimal LES formulations for isotropic turbulence. *Journal of Fluid Mechanics*, 398:321–346, 1999.
- [51] 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, 2021.
- [52] Samuel Lanthaler, Zongyi Li, and Andrew M Stuart. The nonlocal neural operator: Universal approximation. *arXiv preprint arXiv:2304.13221*, 2023.
- [53] Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2009.
- [54] Filippo Santambrogio. Optimal transport for applied mathematicians. *Birkäuser, NY*, 55(58-63):94, 2015.
- [55] Michel Ledoux and Michel Talagrand. *Probability in Banach Spaces: isoperimetry and processes*. Springer Science & Business Media, 2013.
- [56] Nicholas Vakhania, Vazha Tarieladze, and S Chobanyan. *Probability distributions on Banach spaces*, volume 14. Springer Science & Business Media, 2012.

- [57] Michael Brin and Garrett Stuck. *Introduction to dynamical systems*. Cambridge university press, 2002.
- [58] Lan Wen. *Differentiable dynamical systems*, volume 173. American Mathematical Soc., 2016.
- [59] Isaac P Cornfeld, Sergej V Fomin, and Yakov Grigorevich Sinai. *Ergodic theory*, volume 245. Springer Science & Business Media, 2012.
- [60] Roger Temam. *Navier-Stokes equations: theory and numerical analysis*, volume 343. American Mathematical Soc., 2001.
- [61] Qiqi Wang, Rui Hu, and Patrick Blonigan. Least squares shadowing sensitivity analysis of chaotic limit cycle oscillations. *Journal of Computational Physics*, 267:210–224, 2014.
- [62] Sergey P Kuznetsov. *Hyperbolic chaos*. Springer, 2012.
- [63] Stephen Smale. An infinite dimensional version of Sard’s theorem. In *The Collected Papers of Stephen Smale: Volume 2*, pages 529–534. World Scientific, 2000.
- [64] Aly-Khan Kassam and Lloyd N Trefethen. Fourth-order time-stepping for stiff PDEs. *SIAM Journal on Scientific Computing*, 26(4):1214–1233, 2005.
- [65] Gary J Chandler and Rich R Kerswell. Invariant recurrent solutions embedded in a turbulent two-dimensional kolmogorov flow. *Journal of Fluid Mechanics*, 722:554–595, 2013.
- [66] 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 16×16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [67] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [68] J-P Eckmann and David Ruelle. Ergodic theory of chaos and strange attractors. *Reviews of Modern Physics*, 57(3):617, 1985.
- [69] Pritpal Matharu and Bartosz Protas. Optimal closures in a simple model for turbulent flows. *SIAM Journal on Scientific Computing*, 42(1):B250–B272, 2020.

Appendix

In this appendix, we will first provide the supplementary contents of the main text (A), followed by detailed proofs of our theoretical results (B-E), and then present implementation details and more experiment results (F-H). The structure of the appendix is as follows.

- Appendix B provides a list of notations, along with an introduction of important background conceptions, preliminary results, and basic assumptions in this paper.
- Appendix C first formally introduces functional Liouville flow, and then presents a detailed version of Appendix A.5.
- Appendix D provides the proof of the three claims in Theorem 2.1.
- Appendix E provides the proof of Theorem A.1.
- Appendix F contains information about the dataset in the experiments and a visualization of the Navier-Stokes dataset.
- Appendix G provides the implementation details for our method and baseline methods.
- Appendix H first formally introduces the statistics we consider, followed by the full experiment results (table and plots) and ablation studies.

A Supplementary Material

A.1 Comparison Between Different Methods

Table 2: Comparison between different approaches for predicting long-term statistics of Navier-Stokes equations. The Reynolds number Re is large in most applications. The top two are classical approaches, and the rest are machine learning approaches. Training data is counted in the number of snapshots and trajectories. The complexity takes into account both spatial grids and temporal grids. Our approach is even cheaper than coarse-grid simulations because ours can evolve with $O(1)$ time step instead of small time-grids following the CFL condition, as is the case for other methods utilizing a coarse solver. δt is the time-grid size for latent SDE in [38].

Method	Optimal statistics	High-res. training data		Complexity
		FRS Snapshots	Trajs.	
Fully-resolved Simulation, e.g., DNS [43, 44]	✓	-	-	$Re^{3.52}$
Coarse-grid Simulation, e.g., LES [43, 44]	✗	-	-	$Re^{2.48}$
Single-state model [30]	✗	24000	8	$Re^{2.48}$
History-aware model[37]	✗	250000	50	$Re^{2.48}$
Latent Neural SDE[38]	✗	179200	28	$\frac{1}{\delta t} Re^{1.86}$
Online Learning [45]	✗	-	-	$Re^{3.52}$
Physics-Informed Operator Learning (Ours)	✓	110	1	$Re^{1.86}$

A.2 Detailed Background

We formally introduce the problem setting of evaluating long-term statistics as well as existing numerical and machine learning methods. We will show the potential shortcomings of previous learning methods and state some of our theoretical results. See Appendix B for more backgrounds.

A.3 Problem Background

Consider an evolution partial differential equation (PDE) that governs a (nonlinear) dynamical system in the function space,

$$\begin{cases} \partial_t u(x, t) = \mathcal{A}u(x, t) \\ u(x, 0) = u_0(x), u_0 \in \mathcal{H}, \end{cases} \quad (5)$$

where u_0 is the initial value and \mathcal{H} is a function space containing functions of interests, e.g., fluid field, temperature distribution, etc; see Appendix B.4 for further assumptions. This equation naturally

induces a semigroup $\{S(t)\}_{t \geq 0}$ defined as the mapping from the initial state to the state at time t , $S(t) : u_0 \rightarrow u(\cdot, t)$. We refer to the set $\{S(t)u_0\}_{t \geq 0}$ as *trajectory from u_0* .

Attractor, invariant measure, long-term statistics: The (global) *attractor* of the dynamics is defined as the maximal invariant set of $\{S(t)\}_t$ towards which all trajectories converge over time. For many relevant systems, the existence of a compact attractor is either rigorously proved[35, 36] or demonstrated by extensive experiments[46, 47, 48]. The *invariant measure* is the time average of any trajectory, independent of initial value as long as the system is ergodic,

$$\mu^* := \lim_{T \rightarrow \infty} \frac{1}{T} \int_{t=0}^T \delta_{S(t)u} dt, \quad u \in \mathcal{H}, \quad a.e. \quad (6)$$

where δ is the Dirac measure. μ^* is a measure of functions and is supported on the attractor. Intuitively, the invariant measure captures the system’s long-term behavior when it reaches a dynamical equilibrium. The *long-term statistics* are expectations of functionals on the invariant measure. The most straightforward approach to estimate statistics is to run an accurate simulation of trajectory and compute following the definition. In practice, we first fix a sufficiently large T and then choose spatiotemporal grid size accordingly so that the overall error of the simulation within $[0, T]$ remains small. We will refer to this approach as high-fidelity simulations or FRS.

Chaotic systems, characterized by positive Lyapunov exponents [49], are known for their extreme sensitivity to perturbations and catastrophic accumulation of small errors over time. To account for the unstable nature of chaotic systems, high-fidelity simulations have to be carried out on very dense spatiotemporal grids to make discretization errors small enough so that the overall error along the trajectory does not grow rapidly. This makes the FRS approach prohibitively expensive.

A.4 Our Contributions

Our contributions in this work are summarized as follows.

- We propose a novel framework based on functional Liouville flow, to theoretically analyze the problem of estimating long-term statistics of chaotic systems with coarse-grid simulations.
- We formally prove that restricting the learning object in the reduced space, as existing closure models do, suffers from the non-uniqueness of the learning target.
- We leverage physics-informed neural operator as an alternative approach that combines learning on data from both coarse and fine-grid solvers, and physics-based losses. We provide both theoretical and empirical evidence of its superiority.

A.5 Perspective through Liouville Flow in Function Space

We have demonstrated that existing learning methods target a non-unique mapping, resulting in an average of all possible outputs, which can be undesirable. Despite this, these methods still manage to achieve competitive performance. In this section, we show that their empirical result heavily relies on the availability of a large amount of FRS training data. This dependency is a significant limitation, as FRS data are typically scarce. If a sufficient amount of FRS data were already available for training, we could directly compute the statistics using the data, eliminating the need for training a closure model or running coarse-grid simulations.

Our analysis investigates the evolution of the distribution (or measure) to determine whether it converges to μ^* . In terms of finite-dimensional dynamical systems (ODEs), the evolution of distribution is governed by the Liouville equation. This observation motivates us to generalize the Liouville equation into function space and conduct our study therein. Rigorous definitions of related notions and detailed proofs for all claims made in this section can be found in Appendix C.

Functional Liouville Flow: If we expand functions onto an orthonormal basis, $u = \sum_i z_i \psi_i$, a PDE system (of u) can be viewed as an infinite-dimensional ODE (of \mathbf{z}). In this way, we yield the functional version of the Liouville equation describing how the probability density of u evolves. Under this framework, we only need to check the stationary Liouville equation to obtain the limit invariant distribution of a dynamical system and compare it with μ^* .

In the coarse-grid setting, we similarly derive the evolution of the density of \bar{u} and yield the *optimal dynamics* of $v \in \mathcal{F}(\mathcal{H})$ (different from CGS in eq. (2)), v is exactly the same as \bar{u} here),

$$\partial_t v = \mathbb{E}_{u \sim \mu_t} [\mathcal{F} \mathcal{A} u | \mathcal{F} u = v], \quad (7)$$

where μ_t is the distribution of $u \in \mathcal{H}$ following the original dynamics at time t and this expectation is conditioned on the samplings of u satisfying $\mathcal{F} u = v$. Here we arrive at the same result in [50]. Unfortunately, μ_t depends on t and the initial distribution of $u \in \mathcal{H}$, which is underdetermined and will suffer from non-unique issues if restricted to a coarse-grid system, similar to what is discussed in the previous section. In practice, one can only fix one particular $\hat{\mu}$, a distribution in \mathcal{H} , and assign the dynamics in reduced space as $\partial_t v = \mathbb{E}_{u \sim \hat{\mu}} [\mathcal{F} \mathcal{A} u | \mathcal{F} u = v]$. Checking the resulting Liouville equation, we show that μ^* is the choice for $\hat{\mu}$ to guarantee convergence towards $\mathcal{F}_{\#} \rho^*$, the optimal approximation of μ^* in $\mathcal{F}(\mathcal{H})$. Back to the learning methods listed in Section 2, due to the L^2 variational characterization of conditional expectation, the underlying choice of $\hat{\mu}$ is the empirical measure of those training data coming from FRS, ideally μ^* . Consequently, one has to use numerous FRS training data due to the slow convergence of empirical measures for high-dimensional distribution.

As is the case, these learning methods often rely on a large amount of fine-grid data coming from one long FRS trajectory or multiple FRS trajectories which are expensive. Furthermore, most methods still rely on a coarse-grid solver that iteratively evolves with relatively small time steps, and some methods require that the coarse simulation starts from a downsampled version of high-fidelity data close to the attractor. These aspects hinder the further application of these methods.

A.6 Introduction to Operator Learning

The goal of operator learning[31, 32, 51] is to approximate mappings between function spaces rather than vector spaces. One of the representatives is Fourier Neural Operator (FNO) [31], whose architecture can be described as:

$$\mathcal{G}_{FNO} := \mathcal{Q} \circ (W_L + \mathcal{K}_L) \circ \dots \circ \sigma(W_1 + \mathcal{K}_1) \circ \mathcal{P}, \quad (8)$$

where \mathcal{P} and \mathcal{Q} are pointwise lifting and projection operators. The intermediate layers consist of an activation function σ , pointwise operators W_ℓ and integral kernel operators $K_\ell : u \rightarrow \mathcal{F}^{-1}(R_\ell \cdot \mathcal{F}(u))$, where R_ℓ are weighted matrices and \mathcal{F} denotes Fourier transform.

A.7 Provable Convergence to Long-term Statistics:

The universal approximation capability of FNO has been proved[32, 52]. Some might doubt that since small errors will rapidly escalate over time in chaotic systems, we have to train an FNO that perfectly fits the ground truth, which would be unrealistic. However, we have the following result. Intuitively, we show that there exists a true trajectory (from a different initial value) that is consistently close to the simulation we get with approximate FNO. Since the invariant measure is independent of the initial condition, we obtain a good approximation of the (filtered) invariant measure.

Theorem A.1. *For any $h > 0$, denote $\hat{\mu}_{h,\theta} := \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N \delta_{\mathcal{G}_\theta^n v_0(x)}$, any $v_0(x)$ with $x \in D'$. For any $\epsilon > 0$, there exists $\delta > 0$ s.t. as long as $\|(\mathcal{G}_\theta u)(\cdot, h) - S(h)u\|_{\mathcal{H}} < \delta, \forall u \in \mathcal{H}$, we have $\mathcal{W}_{\mathcal{H}}(\hat{\mu}_{h,\theta}, \mathcal{F}_{\#} \mu^*) < \epsilon$, where $\mathcal{W}_{\mathcal{H}}$ is a generalization of Wasserstein distance in function space.*

The proof can be found in Appendix E. Details about $\mathcal{W}_{\mathcal{H}}$ can be found in Appendix B.2. These results show that even if the trained operator jumps a large step h in time and has errors as is in practice, we can still obtain a good estimation of statistics by rolling it out and computing the time average. In practice, a 10% ~ 20% relative error of single-step prediction suffices.

B Notations, Auxiliary Results, and Basic Assumptions

In this section, we first summarize the notations in this paper, then review and define some of the important concepts, and finally state the basic assumptions in this work. We would encourage the readers to always check this section when they have any confusion regarding the proof.

B.1 Notations

Table 3: List of Notations

Notation	Description
μ	Distributions.
$F_{\#}$	Push-forward of a mapping F . If $y = F(x)$ and the distribution of x is μ_x , then the distribution of y is $F_{\#}\mu_x$.
$F^{\#}$	Pull-back of a mapping F . If $y = F(x)$ and the distribution of y is μ_y , then the distribution of x is $F^{\#}\mu_y$.
\mathcal{H}	(Original) Function Space, see eq. (5).
\mathcal{A}	The operator of the dynamics, see eq. (5).
$S(t)$	Semigroup induced by eq. (5).
u	Functions in \mathcal{H} .
\mathcal{F}	Filter. $\mathcal{F} : \mathcal{H} \rightarrow \mathcal{H}$. The image space is finite-rank and denoted as $\mathcal{F}(\mathcal{H})$.
D	The set of grids in fully-resolved simulations. The number of grids is $ D $.
D'	The set of grids in coarse-grid simulations. The number of grids is $ D' $.
$\mathcal{P}(\Omega)$	The set of all probability distributions supported on a set Ω .
$\mathcal{W}_{\mathcal{H}}$	The Wasserstein distance for measures in \mathcal{H} , with $\ \cdot\ _{\mathcal{H}}$ being the cost function.
$\langle \cdot, \cdot \rangle$	The pair of linear functionals and elements in a Banach space \mathcal{X} . Specifically, for $x \in \mathcal{X}$, $f \in \mathcal{X}^*$, its dual space, $\langle f, x \rangle := f(x)$. Thanks to Riesz Representation Theorem, we will also use this notation for inner products in Hilbert space.
\otimes	For a Hilbert space \mathcal{H} , $u \in \mathcal{H}$, $v \in \mathcal{H}$, $u \otimes v$ is defined as the linear operator $w \rightarrow \langle v, w \rangle u$, $w \in \mathcal{H}$.
\oplus	$u \oplus v := (u, v)$.
C_0	Continuous function space, equipped with L^∞ norm.
C_c^∞	Smooth and compactly-supported functions.
$d\mathcal{G}(u, v)$	The Gateaux derivative of operator \mathcal{G} at u in the direction of v .
\aleph_0	Aleph-zero, countably infinite.
$[n]$	$\{1, 2, \dots, n\}$
D_T	The set of time-grids.
$M_{D_T}(u_0)$	The set of functions that are indistinguishable from u_0 merely based on values on spatiotemporal grid $D' \times D_T$, see Assumption D.5.
I_x, I	(Spatial) Grid-measurement operators, see Definition D.1.
$\mathfrak{F}_{D_T}, \mathfrak{F}$	Spatiotemporal grid-measurement operators, see Theorem D.2.
<i>w.r.t.</i>	with regard to

Notation	Description
<i>wlog</i>	without loss of generality

B.2 Optimal Transport in Function Space

Given a Banach space \mathcal{X} and two distributions $\mu_1, \mu_2 \in \mathcal{P}(\mathcal{X})$, we want to measure the closeness of these two distributions.

Recall that in finite-dimensional \mathcal{X} , the (Monge formulation of) c -Wasserstein distance is defined as

$$\mathcal{W}_c(\mu_1, \mu_2) := \inf_T \int_{\mathcal{X}} c(x, Tx) \mu_1(dx), \quad s.t. T_{\#}\mu_1 = \mu_2, \quad (9)$$

where T is a measurable mapping from \mathcal{X} to \mathcal{X} , and $c = c(x, y)$ is non-negative bi-variate function known as cost function.

We could naturally generalize this concept into measures in arbitrary Banach space \mathcal{X} and define the Wasserstein distance correspondingly. In particular, we use the metric in \mathcal{X} as cost function and define

$$\mathcal{W}_{\mathcal{X}}(\mu_1, \mu_2) := \inf_T \int_{\mathcal{X}} \|Tx - x\| \mu_1(dx), \quad s.t. T_{\#}\mu_1 = \mu_2. \quad (10)$$

For more backgrounds and rigorous definitions of concepts appeared above, [53, 54] are standard references for optimal transport, and [55, 56] are good references for measure theory (and probability) in function space.

B.3 Dynamical Systems

We would like to refer the readers to classical textbooks on dynamical systems [57, 58] and ergodic theorems [59] for detailed proofs of lemmas stated in this subsection.

B.3.1 Discrete Time System

Let X be a compact metric space and $f : X \rightarrow X$ be a continuous map. The generic form of discrete-time dynamical system is written as $x_{n+1} = f(x_n)$, $n \in \mathbb{N}$ or $n \in \mathbb{Z}$ (if f is a homeomorphism).

Definition B.1. $\Lambda \subsetneq X$ is an invariant set of f if $f(\Lambda) = \Lambda$.

In the following discussion, we further assume X to be a C^∞ Riemannian manifold without boundary.

Definition B.2. An invariant set $\Lambda \subsetneq X$ of f is hyperbolic if for each $x \in \Lambda$, the tangent space $T_x X$ splits into a direct sum

$$T_x X = E^s(x) \oplus E^u(x), \quad (11)$$

invariant in the sense that

$$Tf(E^s(x)) = E^s(f(x)), \quad Tf(E^u(x)) = E^u(f(x)), \quad (12)$$

such that, for some constant $C \geq 1$ and $\lambda \in (0, 1)$, the following uniform estimates hold:

$$|Tf^n(v)| \leq C\lambda^n |v|, \quad \forall x \in \Lambda, v \in E^s(x), n \geq 0, \quad (13)$$

$$|Tf^n(v)| \geq \frac{1}{C}\lambda^{-n} |v|, \quad \forall x \in \Lambda, v \in E^u(x), n \geq 0. \quad (14)$$

Lemma B.3. (Shadowing Lemma) Let $\Lambda \subset X$ be a hyperbolic set of f . For any $\epsilon > 0$, there is $\eta_0, \eta_1 > 0$ such that for any $\{x_n\}_{n \in \mathbb{N}}$ satisfying (i) $d(x_n, \Lambda) < \eta_0$; (ii) $|x_{n+1} - f(x_n)| < \eta_1$ for all n , there exists $y \in X$ such that $|x_n - f^{(n)}y| < \epsilon$, $\forall n$.

B.3.2 Continuous Time System

Recall that the dynamical system we consider is

$$\begin{cases} \partial_t u(x, t) = \mathcal{A}u(x, t) \\ u(x, 0) = u_0(x), u_0 \in \mathcal{H}, \end{cases} \quad (15)$$

where u_0 is the initial value and \mathcal{H} is a function space containing functions of interests. We will occasionally refer to Au as *vector field governing the dynamics*. This dynamics induced a semigroup $\{S(t)\}_{t \geq 0}$.

Definition B.4. Given a measure $\mu \in \mathcal{P}(\mathcal{H})$, the system is mixing if for any measurable set $A, B \subset \mathcal{H}$, $\lim_{t \rightarrow \infty} \mu(A \cap S(t)(B)) = \mu(A)\mu(B)$.

Definition B.5. A measure $\mu \in \mathcal{P}(\mathcal{H})$ is said to be an invariant measure of this system if $S(t)_\# \mu = \mu$ for all $t > 0$.

Lemma B.6. If a system is mixing, then it is ergodic.

For ergodic systems, there is an invariant measure independent of the initial condition, defined as

$$\mu^* := \lim_{T \rightarrow \infty} \frac{1}{T} \int_{t=0}^T \delta_{S(t)u} dt, \quad u \in \mathcal{H}, \text{ a.e.} \quad (16)$$

where δ is the Dirac measure.

Definition B.7. The semigroup is said to be uniformly compact for t large, if for every bounded set $B \subset \mathcal{H}$ there exists t_0 which may depend on B such that $\cup_{t \geq t_0} S(t)B$ is relatively compact in \mathcal{H} .

Lemma B.8. If $S(t)$ is uniformly compact, then there is a compact attractor in this system.

B.4 Assumptions

Without loss of generality, we carry out our discussion in the regime where \mathcal{H} is a separable Hilbert space to make the proof more readable and concise. We also make the following technical assumptions.

Assumption B.9. \mathcal{H} can be compactly embedded into C_0 .

Assumption B.10. The system eq. (15) is mixing. The semigroup is uniformly compact.

Assumption B.11. The attractor and invariant measure are unique.

Assumption B.12. The attractor is hyperbolic w.r.t $S(t)$ for any $t > 0$.

We remark that these assumptions are either proved or supported by experimental evidence in many real scenarios [60, 61, 46, 62].

For brevity, we will ignore the difference between fully-resolved simulation (FRS) and the exact solution to eq. (15) in the following discussions.

C Formal Introduction of Functional Liouville Flow

In this section, we first formally introduce functional Liouville flow in Appendix C.1. Based on this theoretical framework, we will reformulate the task of estimating long-term statistics of dynamical systems with coarse-grid simulations in Appendix C.2. Finally, we will provide in Appendix C.3 a detailed version of the discussion in Appendix A.5.

C.1 Framework of Functional Liouville Flow for studying Invariant Measure

Functions as vectors: As a corollary of Hahn-Banach theorem, we could always construct a set of orthonormal basis $\{\psi_i\}_i$ of \mathcal{H} such that the filtered space $\mathcal{F}(\mathcal{H}) = \text{span}\{\psi_1, \dots, \psi_n\}$, where we usually have $n = |D'|$. For any function $u \in \mathcal{H}$, there exists a unique decomposition $u(x) = \sum_{i=1}^{\infty} z_i \psi_i(x)$ with $z_i = \langle u, \psi_i \rangle$. This canonically induces an isometric isomorphism:

$$\mathcal{T} : \mathcal{H} \rightarrow \ell^2, \quad u \mapsto \mathbf{z} = (z_1, z_2, \dots). \quad (17)$$

By this means, we can rewrite the original PDE(15) into an ODE in ℓ^2 , denoted by

$$\frac{d\mathbf{z}}{dt} = f(\mathbf{z}), \text{ where } f(\mathbf{z}) \in \ell^2, \quad f(\mathbf{z})_i = \langle \psi_i, \mathcal{A} \circ \mathcal{T}^{-1} \mathbf{z} \rangle, \quad (18)$$

where $\langle \cdot, \cdot \rangle$ is the inner product in \mathcal{H} .

Example C.1. For Kuramoto–Sivashinsky Equation

$$\partial_t u + u \partial_x u + \partial_{xx} u + \partial_{xxxx} u = 0, \quad (x, t) \in [0, 2\pi] \times \mathbb{R}_+, \quad (19)$$

if we choose $\{\psi_k\}$ as the Fourier basis $\{e^{ikx}\}_{k \in \mathbb{Z}}$, then z_k is the coefficient of k -th Fourier mode and the ODE for \mathbf{z} is (component-wise),

$$\frac{dz_k}{dt} = (-k^4 + k^2)z_k - \frac{ik}{2} \sum_{j+l=k} z_j z_l. \quad (20)$$

One could further make z_k real numbers by choosing \sin, \cos basis.

Functional Liouville flow: Recall that in ODE system $\frac{dx}{dt} = f(x)$, $x \in \mathbb{R}^d$, if the initial state x_0 follows the distribution μ_0 whose probability density is $\rho_0(x)$, then the probability density of $x(t)$, denoted by $\rho(x, t)$, satisfies the Liouville equation,

$$\partial_t \rho = -\nabla \cdot (f\rho). \quad (21)$$

Now we want to generalize this result into function space. We need to address the issue that there is in general no probability density function for measures in function space. We will show that it is reasonable to carry on our study by fixing a sufficiently large N and investigating the truncated system of the first N basis, and viewing the densities as the (weak-)limit when $N \rightarrow \infty$.

Proposition C.2. For any μ supported on a bounded set $B \subset \mathcal{H}$ and any $\epsilon > 0$, there exists t_0 and N s.t. for any $u_0 \sim \mu$ and any $t > t_0$, if we write $S(t)u_0$ as $\sum_{i=0}^{\infty} z_i \psi_i$, then $\|\sum_{i>N} z_i \psi_i\| < \epsilon$.

Proof. Define $Q_m := \sum_{i>m} \psi_i \otimes \psi_i$. Then the statement is equivalent to $\|Q_N S(t)u_0\| < \epsilon$, for all $u \in B, t > t_0$.

Due to Assumption B.10, there exists t_0 such that $\cup_{t>t_0} S(t)B$ is relatively compact. This implies that there exists finite (denoted by N_1) points $\{u_i\}$ satisfying that for any $u_0 \in B$, $t > t_0$, there exists $i \leq N_1$ s.t. $\|S(t)u_0 - u_i\| < \frac{\epsilon}{3}$. We define $M_i := \min_j \{j \mid \|Q_j u_i\| < \frac{\epsilon}{5}\}$. We have $M_i < \infty, \forall i$. Choosing N as $\max_{i \leq N_1} M_i$ completes the proof. \square

Remark C.3. We can always restrict our discussion within distributions supported on bounded set whose complement occurs with a probability smaller than machine precision.

Back to the dynamics eq. (15) or the equivalent ODE eq. (18), we first make a generalization. Since the invariant measure is independent of initial condition, we know that for any distribution $\mu_0 \in \mathcal{P}(\mathcal{H})$ (instead of only delta distributions at u_0) from which we sample random initial conditions $u_0 \sim \mu_0$ and evolve these functions, the long-term average $\lim_{T \rightarrow \infty} \frac{1}{T} \int_{t=0}^T (S(t)_\#) \mu_0 dt$ will still converge to μ^* . We will carry out our discussion in this generalized setting where the initial condition is sampled from a distribution μ_0 in function space. We will denote $(S(t)_\#) \mu_0$, the distribution at time t , as μ_t , and denote their density functions for corresponding \mathbf{z} as $\rho(\cdot, t)$ (i.e., $\mu_t = \mathcal{T}_\# \rho(\cdot, t)$). For brevity, we will view $u \in \mathcal{H}$ and $\mathbf{z} \in \ell^2$ as the same and not mention $\mathcal{T}_\#$ or $\mathcal{T}_\#$ for μ_t and $\rho(\cdot, t)$.

If we use the component-form of f , $f = (f_1, f_2, \dots)$, with each f_i a mapping from $\ell^2 \rightarrow \mathbb{R}$, with exactly the same argument to derive eq. (21), we have

$$\partial_t \rho(\mathbf{z}, t) = - \sum_i \partial_{z_i} (f_i(\mathbf{z}) \rho(\mathbf{z}, t)) := -\nabla_{\mathbf{z}} \cdot (f\rho), \quad \rho(\mathbf{z}, 0) = \rho_0(\mathbf{z}). \quad (22)$$

We will refer to this as functional Liouville flow, i.e., the Liouville equation in function space, and denote the R.H.S. operator $\rho \rightarrow -\nabla_{\mathbf{z}} \cdot (f\rho)$ as $\mathcal{L}\rho$.

Reinterpretation of Invariant Measure: With functional Liouville flow, we obtain a new characterization of invariant measure μ^* (whose distribution is denoted as ρ^*).

Proposition C.4. ρ^* is the solution to stationary Liouville equation $\mathcal{L}\rho = 0$.

Proof. Denote $p(\mathbf{z}, t) := \frac{1}{t} \int_{s=0}^t \rho(\mathbf{z}, s) ds$, the finite-time average distribution.

Note that for any \mathbf{z} ,

$$\rho(\mathbf{z}, t) = \rho(\mathbf{z}, 0) + \int_{s=0}^t \partial_t \rho(\mathbf{z}, s) ds \quad (23)$$

$$= \int_0^t \mathcal{L} \rho(\mathbf{z}, s) ds + \rho(\mathbf{z}, 0) \quad (24)$$

$$= \mathcal{L} \int_{s=0}^t \rho(\mathbf{z}, s) ds + \rho(\mathbf{z}, 0) = \mathcal{L}(tp(\mathbf{z}, t)) + \rho(\mathbf{z}, 0). \quad (25)$$

Also, we have $\rho(\mathbf{z}, t) = \partial_t (\int_{s=0}^t \rho(\mathbf{z}, s) ds) = \partial_t (tp(\mathbf{z}, t))$, we conclude that

$$\partial_t (tp(\mathbf{z}, t)) = t \mathcal{L} p(\mathbf{z}, t) + \rho(\mathbf{z}, 0). \quad (26)$$

From this, we yield

$$\partial_t p(\mathbf{z}, t) = \mathcal{L} p(\mathbf{z}, t) + \frac{1}{t} (\rho(\mathbf{z}, 0) - p(\mathbf{z}, t)). \quad (27)$$

By definition we know $p(\mathbf{z}, t) \rightarrow \rho^*$ as $t \rightarrow \infty$, thus $\partial_t p \rightarrow 0$. The term $\frac{1}{t} (\rho(\mathbf{z}, 0) - p(\mathbf{z}, t))$ will also tend to zero as $t \rightarrow \infty$ (recall that they are probability density and thus are uniformly bounded in L^1). Therefore, the limit density ρ^* satisfies $\mathcal{L} \rho^* = 0$. \square

C.2 Reformulation of Estimating Long-term Statistics

We reformulate the problem of *estimating long-term statistics with coarse-grid simulation* with the help of functional Liouville flow. Recall that D' is the set of coarse grid points, and coarse-grid simulation (CGS) is equivalent to evolving functions in $\mathcal{F}(\mathcal{H})$, where \mathcal{F} is the filter (see Section 2).

C.2.1 Notations

We start by defining several notations.

Define the orthonormal projection onto $\mathcal{F}(\mathcal{H})$ as $P = \sum_{i=1}^n \psi_i \otimes \psi_i$. We remark here that in many situations, we have $P = \mathcal{F}$.

Let us decompose \mathbf{z} and u into the resolved part and unresolved part,

$$\mathbf{z} = \mathbf{v} \oplus \mathbf{w}, \quad \mathbf{v} := (c_1, c_2, \dots, c_n), \quad \mathbf{w} := (c_{n+1}, c_{n+2}, \dots); \quad (28)$$

$$u(x) = v(x) + w(x), \quad v(x) := \mathcal{T}^{-1} \mathbf{v} = Pu, \quad w(x) := \mathcal{T}^{-1} \mathbf{w} = (I - P)u. \quad (29)$$

In particular, $w \in \mathcal{F}(\mathcal{H})^\perp$ is the unresolved part in coarse-grid simulations. With this decomposition, we rewrite any density $\rho(\mathbf{z})$ as a joint distribution $\rho(\mathbf{v}, \mathbf{w})$ and define marginal distribution for \mathbf{v} as $\rho_1(\mathbf{v})$, and the conditional distribution of \mathbf{w} given \mathbf{v} as $\rho(\mathbf{w}|\mathbf{v})$. With a little abuse of notation, we will occasionally refer to the probability density as its distribution, and vice versa.

We will also divide the vector field f into resolved part f_r and unresolved part f_u , which are (f_1, f_2, \dots, f_n) and $(f_{n+1}, f_{n+2}, \dots)$ respectively.

C.2.2 Reformulation of Coarse-grid Simulation

We first show that the optimal approximation of μ^* (or ρ^*) in the reduced space is its marginal distribution, if we construct densities with orthonormal basis, as is in eq. (17).

Proposition C.5. $\rho_1^* = \arg \min_{\mu \in \mathcal{P}(\mathcal{F}(\mathcal{H}))} \mathcal{W}_{\mathcal{H}}(\mu, \mu^*)$.

Proof. From the construction of P and the definition of $\mathcal{W}_{\mathcal{H}}$, for any measurable mapping $\mathfrak{T} : \mathcal{H} \rightarrow \mathcal{F}(\mathcal{H})$,

$$\int_{\mathcal{H}} \|\mathfrak{T}u - u\|_{\mu^*}^2(du) \geq \int_{\mathcal{H}} \|Pu - u\|_{\mu^*}^2(du). \quad (30)$$

Thus, for any $\mu \in \mathcal{P}(\mathcal{F}(\mathcal{H}))$,

$$\mathcal{W}_{\mathcal{H}}(\mu, \mu^*) \geq \int_{\mathcal{H}} \|Pu - u\| \mu^*(du) = \mathcal{W}_{\mathcal{H}}(P_{\#}\mu^*, \mu^*). \quad (31)$$

Note that $\rho_1^* = P_{\#}\mu^*$, this completes the proof. \square

This result motivates us to check the evolution of $\rho_1(\mathbf{v}, t)$, which should achieve the optimal approximation ρ_1^* .

Note that by definition, for any distribution $\rho \in \mathcal{P}(\mathcal{H})$, $\rho_1(\mathbf{v}) = \int \rho(\mathbf{v}, \mathbf{w}) d\mathbf{w}$. Combine this with eq. (22), we yield

$$\partial_t \rho_1(\mathbf{v}, t) = \int \partial_t \rho(\mathbf{v}, \mathbf{w}, t) d\mathbf{w} \quad (32)$$

$$= - \int \nabla_{\mathbf{v}} \cdot (f_r(\mathbf{v}, \mathbf{w}) \rho(\mathbf{v}, \mathbf{w}, t)) d\mathbf{w} - \int \nabla_{\mathbf{w}} \cdot (f_u(\mathbf{v}, \mathbf{w}) \rho(\mathbf{v}, \mathbf{w}, t)) d\mathbf{w} \quad (33)$$

$$= - \nabla_{\mathbf{v}} \cdot \left(\frac{\rho_1(\mathbf{v}, t)}{\rho_1(\mathbf{v}, t)} \int f_r(\mathbf{v}, \mathbf{w}) \rho(\mathbf{v}, \mathbf{w}, t) d\mathbf{w} \right) - 0 \quad (34)$$

$$= - \nabla_{\mathbf{v}} \cdot \left(\rho_1(\mathbf{v}, t) \int f_r(\mathbf{v}, \mathbf{w}) \frac{\rho(\mathbf{v}, \mathbf{w}, t)}{\int \rho(\mathbf{v}, \mathbf{w}', t) d\mathbf{w}'} d\mathbf{w} \right) \quad (35)$$

$$= - \nabla_{\mathbf{v}} \cdot (\rho_1(\mathbf{v}, t) \mathbb{E}_{\mathbf{w} \sim \rho(\mathbf{w}|\mathbf{v}; t)} [f_r(\mathbf{v}, \mathbf{w}) | \mathbf{v}]). \quad (36)$$

where we use the divergence theorem for the second term in the second line.

The corresponding ODE dynamics for this Liouville equation eq. (36) is

$$\frac{d\mathbf{v}}{dt} = \mathbb{E}_{\mathbf{w} \sim \rho(\mathbf{w}|\mathbf{v}; t)} [f_r(\mathbf{v}, \mathbf{w}) | \mathbf{v}]. \quad (37)$$

If we transform it back into \mathcal{H} space, it becomes (informally)

$$\partial_t v = \mathbb{E}_{u \sim \mu_t} [\mathcal{F} \mathcal{A} u | \mathcal{F} u = v], \quad (38)$$

as is presented in Appendix A.5 in main text. This describes (one of) the optimal dynamics in the reduced space.

C.2.3 The Effect of Closure Modeling

Apart from the original motivation of closure modeling to approximate the commutator $\mathcal{F} \mathcal{A} - \mathcal{A} \mathcal{F}$, we alternatively interpret it as assigning a vector field \mathcal{A}_{θ} in the reduced space $\mathcal{F}(\mathcal{H})$ and accordingly the coarse-grid dynamics is

$$\partial_t v = \mathcal{A}_{\theta} v, \quad (39)$$

here \mathcal{A}_{θ} plays the role of $\mathcal{A}v + \text{clos}(v; \theta)$ in eq. (2).

We will refer to both \mathcal{A}_{θ} and $\text{clos}(\cdot; \theta)$ as the target of closure modeling for brevity.

As an application of Proposition C.4, we only need to check the solution to the stationary Liouville equation related to this dynamics to decide whether or not the resulting limit distribution is the optimal one ρ_1^* .

C.3 Details for Discussion in Appendix A.5

The dynamics of the filtered trajectory in Equation (38) (we will refer to the equivalent version eq. (37) for convenience), which is also derived in [50], has inspired many works for the design of closure models. Unfortunately, we want to point out that it is impractical to utilize this result for closure model design.

The decision regarding subsequent motion at the state (\mathbf{v}, t) have to made only based on information from the reduced space, which contains merely \mathbf{v} itself and the distribution of \mathbf{v} . For any given \mathbf{v} , only one prediction can be made for the next time step. Similar to the non-uniqueness issue highlighted in Section 2, however, since $\rho(\mathbf{w}|\mathbf{v}; t)$ depends on t and ρ_0 (the initial distribution of $u \in \mathcal{H}$), typically there are multiple distinct $\rho(\mathbf{v}, \mathbf{w}, t)$ with exactly the same \mathbf{v} and marginal distribution in $\mathcal{F}(\mathcal{H})$.

In practice, if one hopes to follow the form of conditional expectation as in eq. (37), he can only fix one particular $q(\mathbf{v}, \mathbf{w})$, a distribution in \mathcal{H} , and assign the vector field in reduced space as $\mathbb{E}_{\mathbf{w} \sim q(\mathbf{w}|\mathbf{v})}[f_r(\mathbf{v}, \mathbf{w})|\mathbf{v}]$.

Now, we check the limit distribution we will obtain with this dynamics. From Proposition C.4, we know that limit distribution $\hat{\rho}_1(\mathbf{v})$ is the solution to (usually in weak sense)

$$\nabla_{\mathbf{v}} \cdot \left(\mathbb{E}_{\mathbf{w} \sim q(\mathbf{w}|\mathbf{v})}[f_r(\mathbf{v}, \mathbf{w})|\mathbf{v}] \rho_1(\mathbf{v}) \right) = 0. \quad (40)$$

Proposition C.6. ρ_1^* is the solution to eq. (40) if $q = \rho^*$.

Proof. By definition, $\rho^*(\mathbf{v}, \mathbf{w})$ satisfies

$$\nabla_{\mathbf{v}} \cdot (f_r(\mathbf{v}, \mathbf{w}) \rho^*(\mathbf{v}, \mathbf{w}, t)) + \nabla_{\mathbf{w}} \cdot (f_u(\mathbf{v}, \mathbf{w}) \rho^*(\mathbf{v}, \mathbf{w}, t)) = 0. \quad (41)$$

Integral over \mathbf{w} and use divergence theorem, we yield

$$0 = \int \nabla_{\mathbf{v}} \cdot (f_r(\mathbf{v}, \mathbf{w}) \rho^*(\mathbf{v}, \mathbf{w})) d\mathbf{w} + 0 \quad (42)$$

$$= \nabla_{\mathbf{v}} \cdot \int f_r(\mathbf{v}, \mathbf{w}) \rho_1^*(\mathbf{v}) \rho^*(\mathbf{w}|\mathbf{v}) d\mathbf{w} \quad (43)$$

$$= \nabla_{\mathbf{v}} \cdot \left(\mathbb{E}_{\mathbf{w} \sim \rho^*(\mathbf{w}|\mathbf{v})}[f_r(\mathbf{v}, \mathbf{w})|\mathbf{v}] \rho_1(\mathbf{v}) \right) \quad (44)$$

This gives the proof. \square

Thus, we show that ρ^* is the correct choice for q to guarantee convergence towards ρ_1^* in $\mathcal{F}(\mathcal{H})$. Back to the learning methods discussed in Section 2, if we follow the new interpretation in Appendix C.2.3, the loss function is

$$J_{ap}(\theta) = \mathbb{E}_{u \sim p_{data}} \|\mathcal{A}_\theta \mathcal{F}u - \mathcal{F} \mathcal{A}u\|^2 \quad (45)$$

$$= \mathbb{E}_{(\mathbf{v}, \mathbf{w}) \sim p_{data}(\mathbf{v}, \mathbf{w})} |f_r(\mathbf{v}; \theta) - f_r(\mathbf{v}, \mathbf{w})|^2, \quad (46)$$

where we transform the original objective function into ℓ^2 space of \mathbf{z} in the second line, and $f_r(\cdot; \theta)$ is the counterpart of \mathcal{A}_θ in ℓ^2 , p_{data} is the empirical measure of training data from fully-resolved simulations(FRS).

Due to the L^2 variational characterization of conditional expectation, the underlying choice of $q(\mathbf{v}, \mathbf{w})$ in those existing learning methods is p_{data} . Consequently, one has to use numerous FRS training data due to the slow convergence of empirical measure of the high-dimensional distribution ρ^* .

D Proof of Theorem 2.1

For the first claim in Theorem 2.1, it has already been shown in the main text that the mapping of closure model $\bar{u} \rightarrow (\mathcal{F} \mathcal{A} - \mathcal{A} \mathcal{F})u$ is not well-defined. We will make this claim more precise in Appendix D.1, and then give the proof for the second claim in Appendix D.2 and the proof for the third claim in Appendix D.3.

D.1 Proof of Theorem 2.1(i)

By transforming the original dynamics into the space of ℓ^2 , it is easier to see why the mapping of closure model is not well-defined. Since $\mathcal{A} \mathcal{F}u = \mathcal{A} \bar{u}$, we only need to show that $\bar{u} \rightarrow \mathcal{F} \mathcal{A}u$ is not well defined. The counterpart of this mapping in ℓ^2 space is $\mathbf{v} \rightarrow f_r(\mathbf{v}, \mathbf{w})$. If it were a well-defined mapping, there would be a mapping $\tilde{f}_r(\mathbf{v})$ such that $f_r(\mathbf{v}, \mathbf{w}) \equiv \tilde{f}_r(\mathbf{v})$ for all \mathbf{w} . In other words, the reduced system is independent of the unresolved part. This property rarely holds in most dynamical systems, except for a few trivial cases like the heat equation.

Next we show that the approximation error has a positive lower bound.

We could always construct $u_1, u_2 \in \mathcal{H}$ such that $\bar{u}_1 = \bar{u}_2$ and $\mathcal{F}\mathcal{A}u_1 \neq \mathcal{F}\mathcal{A}u_2$. Therefore, for any model $\text{clos}(\bar{u}; \theta)$ the approximation error

$$\sup_{u \in \mathcal{H}} \|\text{clos}(\bar{u}; \theta) - (\mathcal{F}\mathcal{A} - \mathcal{A}\mathcal{F})u\|_{\mathcal{H}} \quad (47)$$

$$\geq \sup_{u \in \{u_1, u_2\}} \|\text{clos}(\bar{u}; \theta) - (\mathcal{F}\mathcal{A} - \mathcal{A}\mathcal{F})u\|_{\mathcal{H}} \quad (48)$$

$$\geq \frac{1}{2} (\|(\mathcal{F}\mathcal{A} - \mathcal{A}\mathcal{F})u_1 - \text{clos}(\bar{u}; \theta)\|_{\mathcal{H}} + \|\text{clos}(\bar{u}; \theta) - (\mathcal{F}\mathcal{A} - \mathcal{A}\mathcal{F})u_2\|_{\mathcal{H}}) \quad (49)$$

$$\geq \frac{1}{2} \|(\mathcal{F}\mathcal{A} - \mathcal{A}\mathcal{F})u_1 - (\mathcal{F}\mathcal{A} - \mathcal{A}\mathcal{F})u_2\|_{\mathcal{H}} \quad (50)$$

$$= \frac{1}{2} \|\mathcal{F}(\mathcal{A}u_1 - \mathcal{A}u_2)\|_{\mathcal{H}} \quad (51)$$

has a lower bound independent of the model, where we apply the fact that $\mathcal{F}u_1 = \mathcal{F}u_2 = \bar{u}$ in the last line.

D.2 Proof of Theorem 2.1(ii)

Notations: Recall that \mathcal{H} is the function space, and D' is the set of coarse grid points, $D' = \{x_1, x_2, \dots, x_n\}$. The filtered value of two functions being the same is equivalent to the fact that these two functions have the same values on the grid points in D' .

Definition D.1. Define grid-measurement operator (at x_0)

$$I_{x_0} : \mathcal{H} \rightarrow \mathbb{R} : u \mapsto u(x_0) \quad (52)$$

For brevity, we will use I_j for I_{x_j} . We further define $I_{D'}$ (abbreviated as I if there is no ambiguity),

$$I_{D'} : \mathcal{H} \rightarrow \mathbb{R}^n, u \mapsto (u(x_1), u(x_2), \dots, u(x_n))^T. \quad (53)$$

Before we delve into the details of the proof, we would like to remind the readers of heat equation as an concrete and easy-to-check example where our result holds.

Our original theorem is stated for a continuous time interval. We first prove its finite version.

Theorem D.2. Given D_T the set of time grids, with $|D_T| = N$ and $D_T = \{t_1, t_2, \dots, t_N\}$, for any $r \in \mathbb{N}$, any function $u_0 \in \mathcal{H} \cap C_0$, there exists an r -dimensional manifold $M_r \subset \mathcal{H} \cap C_0$ such that

$$IS(t)u' = IS(t)u_0, \forall t \in D_T, \forall u' \in M_r. \quad (54)$$

Proof. For $m \in \mathbb{N}$, given m linearly independent functions $\{\phi_i\}_{1 \leq i \leq m} \subset \mathcal{H}$, we can construct an affine manifold $A := u_0 + \text{span}\{\phi_1, \dots, \phi_m\}$ and define the following mapping (spatiotemporal grid-measurement operator):

$$\mathfrak{F} = \mathfrak{F}_{D_T} : A \rightarrow \mathbb{R}^{nN} \quad (55)$$

$$v \mapsto \bigoplus_{j=1}^N IS(t_j)v. \quad (56)$$

Note that we have a canonical coordinate system for A :

$$A \leftrightarrow \mathbb{R}^m \quad (57)$$

$$v = u_0 + \sum_{i=1}^m c_i \phi_i \leftrightarrow (c_1, \dots, c_m), \quad (58)$$

thus, \mathfrak{F} is a mapping between finite-dimensional manifolds, and we can compute its Jacobian $\mathcal{J}(v) \in \mathbb{R}^{m \times Nn}$, whose elements consist of the grid-measurement of Gateaux derivatives, $I_j dS(t_k)(v, \phi_l)$, $j \in [n]$, $k \in [N]$, $l \in [m]$.

By a generalization of Sard's theorem for Banach manifold [63], we know that for any u_0 , any r , there exists $m = Nn + r$ linearly independent functions $\{\phi_i\}_{i=1}^m$ such that the Jacobian is everywhere full-rank (i.e., rank = Nn) in the affine manifold A . By pre-image theorem, since $\mathfrak{F}^{-1}\{\mathfrak{F}u_0\}$ is non-empty (for at least $\mathfrak{F}u_0$ is in this set), it is an $m - Nn = r$ dimensional manifold. This gives the proof. \square

Corollary D.3. Given D_T the set of time grids, with $|D_T| = N$ and $D_T = \{t_1, t_2, \dots, t_N\}$, for any $r \in \mathbb{N}$, any function $u_0 \in \mathcal{H} \cap C_0$, there exists infinite $u' \in \mathcal{H}$ such that $IS(t)u' = IS(t)u_0$ for all $t \in D_T$.

Proof. For any $r \geq 1$, there are infinite points in the manifold we yield in the theorem above. \square

Remark D.4. We require $u \in C_0$ only to exclude the artificial case of modifying the function value on a zero-measure set.

From the result above, we see that for any u_0 and finite time-grid set D_T , $\mathfrak{F}^{-1}\{\mathfrak{F}u_0\}$ is an infinite-dimensional manifold. To complete our proof, we make two technical assumptions. One can check these assumptions for specific dynamical systems to derive the final result. We also remark that they are not the weakest set of assumptions to guarantee the final result, we adopt them here primarily to keep the proof concise.

For a time-grid set D_T , and a function u_0 , we denote as $M_{D_T}(u_0)$ the set of all functions u' that $IS(t)u' = IS(t)u_0$, $\forall t \in D_T$.

Assumption D.5. For every u_0 and finite D_T , the infinite-dimensional manifold $M_{D_T}(u_0)$ is unbounded.

Assumption D.6. Given any finite τ , and an arbitrary bounded set $\Omega \subset \mathcal{H}$, we could assign a sequence of linearly independent functions $\{\phi_{i;u}\}_{i=1}^\infty \subset \mathcal{H}$ for each $u \in \mathcal{H}$ such that for any functions u, v , any subset $B \subset \mathbb{N}$, and any finite subset D_T of $[0, \tau]$, there exists a continuous mapping

$$G : M_{D_T}(u) \cap \{u + \text{span}\{\phi_{i;u} \mid i \in B\}\} \rightarrow M_{D_T}(v) \cap \{v + \text{span}\{\phi_{i;v} \mid i \in B\}\} \quad (59)$$

depending only on u, v, B, D_T , such that

$$\sup_{w \in \Omega \cap M_{D_T}(u)} \|Gw - w\| \leq C_\Omega \|u - v\|, \quad (60)$$

where the constant C_Ω only depends on Ω and τ , and not on u, v, B, D_T .

Before moving on, we first review a classical result.

Lemma D.7. There exists a bijection between \mathbb{N} and \mathbb{N}^2 .

Proof. This is a standard result in set theory, and its proof can be found in many textbooks.

We give an example on how to construct such a bijection (see Figure 2). We write 1,2,3,4,... zigzaggingly to fill the \mathbb{N}^2 plane. In this way, we construct a mapping $\iota : \mathbb{N}^2 \rightarrow \mathbb{N}$, with $\iota(i, j)$ defined as the value written at (i, j) -position in the \mathbb{N}^2 plane. Clearly, ι is a bijection. \square

With ι , we can partition \mathbb{N} into \aleph_0 (countably infinite) disjoint subsequences, $\{\iota(i, j)\}_{j \in \mathbb{N}}$ for each i .




Figure 2: Illustration of a bijection between \mathbb{N} and \mathbb{N}^2 using a zigzag numbering scheme.

We are finally ready for the proof of Theorem 2.1(ii).

Theorem D.8. For any $u \in \mathcal{H}$ and finite τ , there exist *infinite* $u' \in \mathcal{H}$ such that $\mathcal{F}S(t)u' = \mathcal{F}S(t)u$ for all $t \in [0, \tau)$.

Proof. Let \mathcal{K} be a dense Hilbert subspace of \mathcal{H} that can compactly embedded into \mathcal{H} , with norm $\|\cdot\|_{\mathcal{K}}$. For instance, if $\mathcal{H} = H^k$, the Sobolev space $W^{2,k}$, then we can choose \mathcal{K} as H^{k+1} .

Step 1: We first deal with the case when $u \in \mathcal{K}$.

Define a sequence of time-grid set D_j as $\{\frac{i}{2^j}\tau | 0 \leq i < 2^j\}$. Similar to the argument in Theorem D.2, we can construct a sequence $\{\phi_i\}_{i=1}^{\infty} \subset \mathcal{K}$ satisfying the following properties.

- (i) They are linearly independent.
- (ii) For any $j \in \mathbb{N}$, $m > 2^j$, $B \subset \mathbb{N}$ with $|B| = m$, the spatiotemporal grid-measurement \mathfrak{F} for $D' \times D_j$ has full-rank Jacobian everywhere in the affine manifold $u + \text{span}\{\phi_k : k \in B\}$.

Based on the $\mathbb{N} - \mathbb{N}^2$ bijection ι , we define the following subspaces:

$$E_k := \text{span}\{\phi_{\iota(k,i)} : i \in \mathbb{N}\}. \quad (61)$$

There are \aleph_0 such subspaces in total, we will next find a point u' in each E_k such that $IS(t)u' = IS(t)u$ for all $t \in [0, \tau)$.

WLOG, we will only show how to construct u' in E_1 .

We denote

$$M_j = E_1 \cap M_{D_j}(u). \quad (62)$$

By the construction of D_j we have $M_1 \supset M_2 \supset M_3 \supset \dots$

We first fix three constants $0 < B_0 < B_1$, $B_2 > 0$ and and construct a sequence of $\{u_i\} \subset E_1$ such that

- (i) $u_i \in M_i$.
- (ii) $B_0 < \|u_i - u\|_{\mathcal{H}} < B_1$.
- (iii) $\|u_i\|_{\mathcal{K}} < B_2$.

This construction is achievable due to Assumption D.5 and the fact that $u \in \mathcal{K}$.

Since \mathcal{K} can be compactly embedded into \mathcal{H} , there there exists a subsequence $\{u_{i_j}\}$ of $\{u_i\}$ that is convergent in \mathcal{H} . We denote its limit as u_{∞} . From (ii), we have $\|u_{\infty}\|_{\mathcal{H}} < \infty$ and $u_{\infty} \neq u$.

Due to Assumption B.9, we have $u_{i_j} \rightarrow u_{\infty}$ in C_0 , which implies that for any D_j , $\mathfrak{F}_{D_j}u_{\infty} = \mathfrak{F}_{D_j}u$.

Because of the continuity of the mapping $t \mapsto I_x S(t)v$ for any $x \in D'$, $v \in \mathcal{H}$, we know that $IS(t)u' = IS(t)u$ for all $t \in [0, \tau)$.

To conclude, for each E_k , there exists $u' \in E_k$ that is not distinguishable from u merely based on function values restricted to the grid $D' \times D_T$. Recall that for any $j \neq k$, by construction we have $E_j \cap E_k = \{u\}$, thus these u' in different E_k are mutually different. This completes the proof for the case $u \in \mathcal{K}$.

Step 2 Now we give the proof for general $u \in \mathcal{H}$.

Since \mathcal{K} is dense in \mathcal{H} , there exists a sequence $\{u^n\} \subset \mathcal{K}$ such that $\|u^n - u\|_{\mathcal{H}} < \frac{1}{2^n}$. We keep using the constant B_0, B_1, B_2 and define $\Omega := \{v \in \mathcal{H} | \|v - u\|_{\mathcal{H}} < B_1 + 1\}$. Following Assumption D.6, we obtain the constant C_{Ω} and linearly independent set $\{\phi_{i,u^n}\}_{i=1}^{\infty}$ for each n . Following the first part of this proof, we define

$$E_k^n := \text{span}\{\phi_{\iota(k,i);u^n} | i \in \mathbb{N}\} \quad M_{j,k}^n := E_k^n \cap M_{D_j}(u^n), \quad (63)$$

and again we only need to consider the case when $k = 1$, and thus abbreviate $M_{j,k}^n$ as M_j^n .

We will restrict our discussion within $n > n_0 := \left\lceil \max\{\log_2 \frac{6(C_{\Omega}+1)}{B_0}, \log_2 3(C_{\Omega} + 1)\} \right\rceil + 1$.

We inductively construct a sequence (indexed by n) of sequence $\{u_j^n\}_j \subset E_1^n$ as follows:

(I) For $n = n_0$, we construct $\{u_j^n\}_j$ the same as the first part of the proof.

- (i) $u_j^n \in M_j^n$.
- (ii) $B_0 < \|u_j^n - u^n\|_{\mathcal{H}} < B_1$.
- (iii) $\|u_j^n\|_{\mathcal{K}} < B_2$.

(II) Now suppose we have constructed $\{u_j^n\}_j$, we apply Assumption D.6 for u^n, u^{n+1} , $B = \{\iota(1, i) \mid i \in \mathbb{N}\}$ and D_j and obtain a continuous mapping G . We choose u_j^{n+1} as $G u_j^n$.

Next, we give some estimations for u_j^{n+1} .

First, note that we have

$$\|u^n - u^{n+1}\|_{\mathcal{H}} \leq \|u^n - u\|_{\mathcal{H}} + \|u - u^{n+1}\|_{\mathcal{H}} < \frac{3}{2^{n+1}}, \quad (64)$$

and thus $\|u_j^n - u_j^{n+1}\|_{\mathcal{H}} \leq C_{\Omega} \frac{3}{2^{n+1}}$ by construction. Based on this, we have

$$\|u_j^{n+1} - u^{n+1}\|_{\mathcal{H}} \geq \|u_j^n - u^n\|_{\mathcal{H}} - \|u_j^n - u_j^{n+1}\|_{\mathcal{H}} - \|u^n - u^{n+1}\|_{\mathcal{H}} \quad (65)$$

$$\geq \|u_j^n - u^n\|_{\mathcal{H}} - \frac{3(C_{\Omega} + 1)}{2^{n+1}}. \quad (66)$$

By induction, we have

$$\|u_j^{n+1} - u^{n+1}\|_{\mathcal{H}} \geq \|u_j^{n_0} - u^{n_0}\|_{\mathcal{H}} - \sum_{n > n_0} \frac{3(C_{\Omega} + 1)}{2^{n+1}} > \frac{B_0}{2}. \quad (67)$$

We also have

$$\|u_j^{n+1} - u^{n+1}\|_{\mathcal{H}} \leq \|u_j^n - u^n\|_{\mathcal{H}} + \|u_j^n - u_j^{n+1}\|_{\mathcal{H}} + \|u^n - u^{n+1}\|_{\mathcal{H}} \quad (68)$$

$$\leq \|u_j^n - u^n\|_{\mathcal{H}} + \frac{3(C_{\Omega} + 1)}{2^{n+1}}. \quad (69)$$

By induction, we have

$$\|u_j^{n+1} - u^{n+1}\|_{\mathcal{H}} \leq \|u_j^{n_0} - u^{n_0}\|_{\mathcal{H}} + \sum_{n > n_0} \frac{3(C_{\Omega} + 1)}{2^{n+1}} < B_1 + \frac{1}{2}. \quad (70)$$

Similar to what is done in the first part of the proof, we can choose v^{n_0} as one of the limit points of $\{u_j^{n_0}\}_j$. Inductively, we can construct a sequence of v^n such that

- (i) v^n is one of the limit points of $\{u_j^n\}_j$
- (ii) $\|v^n - v^{n-1}\|_{\mathcal{H}} \leq \frac{3C_{\Omega} + 1}{2^n}$.

Thus, $\{v^n\}_n$ is a Cauchy sequence in \mathcal{H} and we denote its limit as v . Because of Assumption B.9, we have that $IS(t)v = IS(t)u, \forall t \in [0, \tau)$. It is also clear that

$$v \in M_{[0, \tau)}(u) \cap \{u + \text{span}\{\phi_{\iota(1, i); u} \mid i \in \mathbb{N}\}\}, \quad (71)$$

and

$$\|v - u\|_{\mathcal{H}} = \left\| \lim_{n \rightarrow \infty} (v^n - u^n) \right\|_{\mathcal{H}} \geq \liminf_{n \rightarrow \infty} \|u^n - v^n\|_{\mathcal{H}} \geq \frac{B_0}{2}. \quad (72)$$

With exactly the same argument as in the first step, we construct infinite mutually-different functions that are not distinguishable from u on $D' \times [0, \tau)$. This completes the proof. \square

D.3 Proof of Theorem 2.1(iii)

Theorem D.9. *One cannot obtain the ρ_1^* if there is randomness in the evolution of dynamics.*

Proof. Our proof will be carried out for a more general setting.

Consider two dynamics

$$\frac{d\mathbf{v}}{dt} = b_1(\mathbf{v}) \quad (73)$$

$$d\mathbf{v} = b_2(\mathbf{v})dt + \sigma(\mathbf{v})dW. \quad (74)$$

The first one corresponds to the deterministic motion as is in the original dynamical system (transformed into ℓ^2). Either ρ^* or ρ_1^* is the limit distribution of certain deterministic dynamics. The second one corresponds to the dynamics of the stochastic closure model. Here dW is a d dimensional Brownian motion where d is the latent dimension of the model.

From Appendix C, we know that the limit distribution of eq. (73) is the solution to stationary Liouville equation,

$$\nabla \cdot (b_1 \rho) = 0. \quad (75)$$

As for eq. (74), similar to how we handle deterministic systems in Appendix C and how we derive the Fokker-Planck equation in finite-dimensional systems, we can generalize Fokker-Planck equation into function space and yield that the limit distribution of eq. (74) is the solution to stationary Fokker-Planck equation,

$$\nabla \cdot (b_2 \rho) - \frac{1}{2} \nabla^2 : (\sigma \sigma^T \rho) = 0. \quad (76)$$

Next, we argue by contradiction to show that ρ_1^* or ρ^* will not satisfy eq. (76).

Suppose ρ_1^* is the solution to both eq. (75) and eq. (76). We then have that for any $k \in \mathbb{R}$, ρ_1^* is the solution to

$$\nabla \cdot (kb_1 + b_2 \rho) - \frac{1}{2} \nabla^2 : (\sigma \sigma^T \rho) = 0. \quad (77)$$

We can expand this equation and write it in the following form

$$A(\mathbf{v}) : \nabla^2 \rho + B_k(\mathbf{v}) \cdot \nabla \rho + C_k(\mathbf{v}) \rho = 0, \quad (78)$$

in particular, $A = -\frac{1}{2} \sigma \sigma^T$ and $C_k(\mathbf{v})$ has the form $C(\mathbf{v}) + k \nabla \cdot b_1$.

Recall that ρ_1^* is supported on the compact attractor (denoted as Ω) of the original dynamical system. Now let us consider the maximum point of ρ_1^* in Ω . If there exists local maximum point $\mathbf{v}_0 \in \Omega^\circ$ (its interior), then the Hessian of ρ_1^* at \mathbf{v}_0 is semi-negative-definite and $\nabla \rho_1^*(\mathbf{v}_0) = 0$. Thus we have

$$C_k(\mathbf{v}_0) \rho_1^*(\mathbf{v}_0) = -A(\mathbf{v}) : \nabla^2 \rho_1^*(\mathbf{v}_0) \geq 0. \quad (79)$$

We could always choose k such that the L.H.S. has a negative value. Contradiction!

This suggests that there is no local maxima of ρ_1^* in Ω° . Thus, the maxima of ρ_1^* is on the boundary. However, $\rho_1^*|_{\partial\Omega} = 0$. This suggests $\rho_1^* = 0$. Contradiction!

This completes the proof. □

E Proof of Theorem A.1

As a preliminary result, we show the following properties of the dynamical systems under consideration.

Lemma E.1. *For any $h > 0$, any initial condition $u_0 \in \mathcal{H}$, $\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N \delta_{S(nh)u_0} = \mu^*$.*

Proof. Denote $G := S(h)$. From Assumption B.10, we know that for any measurable set $A, B \subset \mathcal{H}$, $\lim_{t \rightarrow \infty} \mu(A \cap S(t)(B)) = \mu(A)\mu(B)$. In particular $\lim_{n \rightarrow \infty} \mu(A \cap S(nh)(B)) = \mu(A)\mu(B)$. This suggests that the system defined by

$$u_{n+1} = G u_n \quad (80)$$

is mixing.

From Lemma B.6, this system is ergodic, thus having an invariant measure w.r.t G . Since $G_{\#}\mu^* = \mu^*$, due to the uniqueness of invariant measure, we derive the proof. \square

In the following proof, we will use the notation $G := S(h)$, which is the learning target (ground-truth operator), and \hat{G} for the approximate operator we obtain after training. By the design of the neural operator, the input of \hat{G} can be vectors in \mathbb{R}^d for any dimensionality d , serving as various discretizations of a particular function from \mathcal{H} . Here we violate the concepts a little bit by denoting \hat{G} as a mapping from \mathcal{H} to \mathcal{H} (approximating $S(h)$) instead of \mathcal{H} to $\mathcal{H} \times [0, h]$ (approximating $u \rightarrow \{S(t)u\}_{t \leq h}$) as is in our algorithm in main. In practice, we only use the last element of the output sequence, corresponding to the prediction for $S(h)u$, for estimating statistics. To be specific, to estimate long-term statistics in coarse-grid systems with learned operator \hat{G} , we use as input a function in reduced space $v(x) \in \mathcal{F}(\mathcal{H})$, $x \in D'$ (equivalent to an $\mathbb{R}^{|D'|}$ vector consisting of the function values on the grids), and autoregressively compute $\hat{G}^{(n)}v$, $n \in \mathbb{N}$. The invariant measure is estimated by

$$\hat{\mu}_{D'} := \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N \delta_{\hat{G}^{(n)}v}. \quad (81)$$

We first remind readers of the following fact.

Fact E.2. For any function $u_0 \in \mathcal{H}$, let $\tilde{u} := (u_0(x_1), \dots, u_0(x_n))^T$, $n = |D'|$, be the discretization of u_0 in the coarse-grid system. There exists $u \in \mathcal{H}$ (possibly different from u_0) such that

$$\tilde{u} = I_{D'}u', \quad \hat{G}(\tilde{u}) = I_{D'}\hat{G}u. \quad (82)$$

Now, we prove our main result.

Theorem E.3. For any $h > 0$ and any $\epsilon > 0$, there exists $\delta > 0$ such that, as long as $\|\hat{G}u - Gu\|_{\mathcal{H}} < \delta$, $\forall u \in \mathcal{H}$, we will have $\mathcal{W}_{\mathcal{H}}(\hat{\mu}_{D'}, \rho_1^*) < \epsilon$.

Proof. We will first deal with dynamics in the original space \mathcal{H} . We have two dynamics, the exact one and the approximate one,

$$u^{n+1} = Gu^n, \quad u^0 = u_0 \in \mathcal{H}; \quad (83)$$

$$\hat{u}^{n+1} = \hat{G}\hat{u}^n, \quad \hat{u}^0 = u_0 \in \mathcal{H}. \quad (84)$$

Both dynamics will converge to an attractor, Ω and $\hat{\Omega}$, respectively.

With Lemma B.3, we know that there exists $\eta_0, \eta_1 > 0$ such that for any $\{u_n\}_{n \in \mathbb{N}} \subset \mathcal{H}$ satisfying (i) $d(u_n, \Omega) < \eta_0$; (ii) $\|u_{n+1} - G(u_n)\| < \eta_1$ for all n , there exists $\tilde{u} \in \mathcal{H}$ such that

$$\|u_n - G^{(n)}\tilde{u}\| < \epsilon, \quad \forall n \in \mathbb{N}. \quad (85)$$

From Theorem 1.2 in Chapter 1 of [60], we know that there exists $\eta_2 > 0$ such that

$$\|Gu - \hat{G}u\| < \eta_2, \quad \forall u \in \mathcal{H} \Rightarrow \text{dist}(\Omega, \hat{\Omega}) < \frac{\eta_0}{5}. \quad (86)$$

Now we choose δ as $\min\{\eta_1, \eta_2\}$ and define the approximate operator \hat{G} as well as dynamics and attractor accordingly.

We next choose $n_0 \in \mathbb{N}$ such that $\sup_{n \geq n_0} \text{dist}(\hat{u}^n, \hat{\Omega}) < \frac{\eta_0}{5}$. This implies that $\sup_{n \geq n_0} \text{dist}(\hat{u}^n, \Omega) < \frac{3\eta_0}{5}$. We apply Lemma B.3 to obtain a function $\tilde{u} \in \mathcal{H}$ such that

$$\|u^n - G^{(n-n_0)}\tilde{u}\| < \epsilon, \quad \forall n \geq n_0. \quad (87)$$

For any $N \in \mathbb{N}$, we define

$$\mu_N := \frac{1}{N} \sum_{n=0}^N \delta_{G^{(n)}\tilde{u}} \quad (88)$$

$$\hat{\mu}_N := \frac{1}{N} \sum_{n=n_0}^{n_0+N} \delta_{\hat{u}^n}. \quad (89)$$

By constructing the transport mapping $T : u^n \mapsto G^{(n-n_0)}\tilde{u}$, $n_0 \leq n \leq n_0 + N$, we have that

$$\mathcal{W}_{\mathcal{H}}(\hat{\mu}_N, \mu_N) < \epsilon. \quad (90)$$

Note that $\hat{\mu}_N \rightarrow \hat{\mu}_D$ (estimated invariant measure with fine-grid simulations) as $N \rightarrow \infty$, we derive

$$\mathcal{W}_{\mathcal{H}}(\hat{\mu}_D, \mu^*) \leq \epsilon. \quad (91)$$

Recall that P is the orthonormal projection towards $\mathcal{F}(\mathcal{H})$ and that $\|Pu - Pu'\| \leq \|u - u'\|$ for any $u, u' \in \mathcal{H}$. In light of Fact E.2, we derive $\mathcal{W}_{\mathcal{H}}(\hat{\mu}_D, \rho_1^*) \leq \epsilon$. \square

F Experiment Setup and Data Generation

F.1 Kuramoto–Sivashinsky Equation

We consider the following one-dimensional KS equation for $u(x, t)$,

$$\partial_t u + u\partial_x u + \partial_{xx} u + \nu\partial_{xxx} u = 0, \quad (x, t) \in [0, L] \times \mathbb{R}_+, \quad (92)$$

with periodic boundary conditions. The positive viscosity coefficient ν reflects the traceability of this equation. The smaller ν is, the more chaotic the system is. We study the case for $\nu = 0.01$, $L = 6\pi$.

FRS is conducted with exponential time difference 4-order Runge-Kutta (ETDRK4)[64] with 1024 uniform spatial grid and 10^{-4} time grid. The CGS is conducted with the same algorithm except with 128 uniform spatial grids and 10^{-3} timegrid. We choose $h = 0.1$ for our model.

Dataset The training dataset for the neural operator consists of two parts, the CGS data and FRS data. The CGS dataset contains 6000 snapshots from 100 CGS trajectories. Snapshots are collected from time $t = 20 + k$, $k = 1, 2, \dots, 60$. The data appears as input-label pairs $(v(\cdot, t), v(\cdot, t + h))$, where $h = 0.1$ for KS. The FRS dataset contains 105 snapshots from 3 FRS trajectories. Snapshots are collected from $t = 20 + 2k$, $k = 1, 2, \dots, 35$. The data appears as input-label pairs $(u(\cdot, t), \{u(\cdot, t + \frac{k}{4}h) | k = 1, 2, 3, 4\})$.

As for input functions of PDE loss, they come from adding Gaussian random noise to FRS data.

Estimating Statistics For all methods in this experiment, statistics are computed by averaging over $t \in [20, 150]$ and 400 trajectories with random initializations.

F.2 Navier-Stokes Equation

We consider two-dimensional Kolmogorov flow (a form of the Navier-Stokes equations) for a viscous incompressible fluid (fluid field) $\mathbf{u}(x, y, t) \in \mathbb{R}^2$,

$$\partial_t \mathbf{u} = -(\mathbf{u} \cdot \nabla) \mathbf{u} - \nabla p + \frac{1}{Re} \Delta \mathbf{u} + (\sin(4y), 0)^T, \quad \nabla \cdot \mathbf{u} = 0, \quad (x, y, t) \in [0, L]^2 \times \mathbb{R}_+, \quad (93)$$

with periodic boundary conditions. In the experiment, we deal with the vorticity form of this equation.

$$\partial_t w = -\mathbf{u} \cdot \nabla w + \frac{1}{Re} \Delta w + \nabla \times (\sin(4y), 0)^T, \quad (94)$$

where $w = \nabla \times \mathbf{u}$. The positive coefficient Re is the Reynolds number. The larger Re is, the more chaotic the system is. We consider the case $Re = 100$, $L = 2\pi$.

FRS is conducted with pseudo-spectral split-step [65] with $128 * 128$ uniform spatial grid and self-adaptive time grid. The CGS is conducted with the same algorithm except with $16 * 16$ uniform spatial grids. For our model, we choose $h = 1$.

Dataset The training dataset for the neural operator consists of two parts, the CGS data and FRS data. The CGS dataset contains 8000 snapshots from 80 CGS trajectories. Snapshots are collected from time $t = 80 + 4k$, $k = 1, 2, \dots, 100$. The data appears as input-label pairs $(u(\cdot, t), \{u(\cdot, t + \frac{k}{16}h) | k \in [16]\})$, where $h = 0.1$ for KS. The FRS dataset contains 110 snapshots from 1 FRS trajectories. Snapshots are collected from $t = 50 + 3k$, $k = 1, 2, \dots, 110$. The data appears as input-label pairs $(u(\cdot, t), \{u(\cdot, t + \frac{k}{16}h) | k \in [16]\})$.

As for input functions of PDE loss, they come from adding Gaussian random noise to FRS data.

Estimating Statistics For all methods in this experiment, statistics are computed by averaging over $t \in [1800, 3000]$ and 400 trajectories with random initializations.

G Implementation Details

G.1 Physics-Informed Operator Learning

Algorithm 1 Multi-stage Physics-Informed Operator Learning

Input: Neural operator \mathcal{G}_θ ; training data set $\mathcal{D}_c(\text{CGS})$, $\mathcal{D}_f(\text{FRS})$, \mathcal{D}_p (randomly sampled).
Hyper-parameters: Training iterations N_i ($i = 1, 2, 3$). Weights combining two loss $\lambda_i(t)$ ($i = 1, 2$), which decay as t increases. Parameters regarding optimizer.

- 1: **for** $t = 1, \dots, N_1$ **do**
 - 2: Minimize $J(\theta; \mathcal{D}_c)$
 - 3: **for** $t = 1, \dots, N_2$ **do**
 - 4: Minimize $\lambda_1(t)J_{data}(\theta; \mathcal{D}_c) + J_{data}(\theta; \mathcal{D}_f)$
 - 5: **for** $t = 1, \dots, N_3$ **do**
 - 6: Minimize $\lambda_2(t)J_{data}(\theta; \mathcal{D}_f) + J_{pde}(\theta; \mathcal{D}_p)$
 - 7: **return** \mathcal{G}_θ
-

Following the notations in the main, we formally summarize our algorithm as in *Algorithm 1*.

For input initial value u_0 (function restricted on the grid, which is a 1D tensor for KS and 2D tensor for NS), we repeat u_0 T times to make it a 2D tensor or 3D tensor (u_0, u_0, \dots, u_0) , respectively, where T is a hyperparameter. For the implementations, the neural operator will learn to predict the mapping

$$(u_0, u_0, \dots, u_0) \rightarrow \bigoplus_{j=1}^T S\left(\frac{j}{T}h\right) u_0, \quad (95)$$

which is a discretization of $\{S(t)u_0\}_{t \in [0, h]}$.

KS Equation Following the architecture in the original FNO paper[31], our model is a 4-layer FNO with 32 hidden channels and 64 projection channels. We choose $h = 0.1$, $T = 64$. The data loss will only be computed for the time grid where there is label information.

We first train the model with CGS data, we use ADAM for optimization, with learning rate 5e-2, scheduler gamma 0.7 and scheduler stepsize 100. We train with batchsize 32 for 1000 epochs.

Then we train the model with CGS data and FRS data. $\lambda_1(0) = 1$ and halves every 100 epochs. We train with batchsize 32 for 250 epochs.

Finally, we train the model with PDE loss. We train with batch size 8 for 1487 epochs. Each batch contains 4 functions for computing the data loss and 4 functions for computing the PDE loss. $\lambda_2(t)$ decreases by 1.7 for every 500 epochs.

When we finish training, the L^2 relative error on the FRS test set is $\sim 12\%$.

NS Equation Our model is a 4-layer FNO, with 32 hidden channels and 64 projection channels. We choose $h = 1$, $T = 32$. The data loss will only be computed for time grid where there is label information.

We first train the model with CGS data, we use ADAM for optimization, with learning rate 4e-3, scheduler gamma 0.6 and scheduler stepsize 50. We train with batch size 32 for 60 epochs.

Then we train the model with CGS data and FRS data. $\lambda_1(t) = \mathbf{1}_{t \leq 20}$ and halves every 100 epochs. We train with batch size 8 for 53 epochs.

Finally, we train the model with PDE loss. We train with batch size 16 for 1530 epochs. Each batch contains 8 functions for computing data loss and 8 functions for computing PDE loss. $\lambda_2(t)$ decreases by 1.8 for every 60 epochs.

When we finish training, the L^2 relative error on the FRS test set is $\sim 19\%$. The training takes ~ 40 minutes to complete.

G.2 Baseline Method: Single State Closure Model

The network follows the Vision Transformer [66] architecture. For KS equation, the input was partitioned into 1×4 patches, with 2 transformer layers of 6 heads. The hidden dimension is 96 and the MLP dimension is 128. For NS equation, the input was partitioned into 4×4 patches, with 2 transformer layers of 6 heads. The hidden dimension is 96 and the MLP dimension is 128. For both experiments, we use AdamW optimizer [67] with learning rate $1e-4$ and weight decay $1e-4$.

H More Experiment Results and Visualizations

H.1 Statistics

We formally introduce the statistics we consider.

Total Variation for Invariant Measures As is mentioned in the main text, we propose to directly compare the estimated invariant measure resulting from the time average of simulations and that of ground truth.

Recall that we have expanded $u \in \mathcal{H}$ onto orthonormal basis $u = \sum_{i=1}^{\infty} z_i \psi_i$. In particular, for $v \in \mathcal{F}(\mathcal{H})$, $v \in \text{span}\{\psi_i : i \leq |D'|\}$. We compute the total variation(TV) distance for the (marginal) distribution of each v_i , where TV distance of two distributions (probability densities) ν, μ is defined as

$$d_{TV}(\mu, \nu) = \frac{1}{2} \int |\mu(x) - \nu(x)| dx. \quad (96)$$

For experiments we consider in this work, a natural choice of ψ_i is the Fourier basis functions, $\{e^{i \frac{2k\pi}{L} x}\}_{k \in \mathbb{Z}}$ for 1D KS and $\{e^{i \frac{2\pi}{L} (kx+ jy)}\}_{k, j \in \mathbb{Z}^2}$ for 2D NS.

With a little abuse of definition, the corresponding z_i are complex numbers. [68] shows that the limit distribution of $\text{Arg} z_i$ is uniform distribution on $[0, 2\pi]$. Thus, it suffices to check the distribution of mode length $|z_i|$.

Other Statistics In the following, we use \hat{u}_k to denote the k -th Fourier mode of u . When u is a multi-variate function, k is a tuple.

- Energy Spectrum.
 $O_e(u; k) = |\hat{u}_k|^2$ (1D), $O_e(u; k_0) = \sum_{|k|_1 = k_0} |\hat{u}_k|^2$ (general).
The k_0 -th energy spectrum is $\mathcal{O}_e(k) := \mathbb{E}_{u \sim \mu^*} O_e(u; k_0)$.
- Spatial Correlation.
 $O_s(u; h) = \int u(x)u(x+h)dx$. The h spatial correlation is $\mathcal{O}_s(h) := \mathbb{E}_{u \sim \mu^*} O_s(u; h)$.
- Auto Correlation Coefficient.
Note that \mathcal{O}_s is a function of h .
The k -th Auto Correlation Coefficient is $\mathcal{O}_a(k) := |(\hat{\mathcal{O}}_s)_k|^2$.
- The distribution of Vorticity ($w(x)$ for NS) and Velocity ($u(x)$ for KS).
- The variance of the function value.
- Dissipation Rate: $\frac{1}{Re} \int u(x)^2 dx$, where \int refers to averaged integral

$$\int_{\Omega} f(x) dx := \frac{\int_{\Omega} f(x) dx}{\int_{\Omega} dx}. \quad (97)$$

In practice, we usually check the distribution of this quantity.

- Kinetic Energy: $\int (u - \bar{u})^2 dx$ where $\bar{u}(x) := \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T u(x, t) dt$. In practice, we usually check the distribution of this quantity.

Table 4: **Error on Different Statistics: KS equation** Header: From left to right: Average relative error on energy spectrum, max relative error on energy spectrum, average relative error on auto-correlation coefficient, max error on auto-correlation coefficient, total variation distance from (ground truth) velocity distribution, average component-wise TV distance(error), and max component-wise TV distance(error).

Method	Avg. Eng.	Max Eng.	Avg. Cor.	Max Cor.	Velocity	Avg. TV	Max TV
CGS (No closure)	12.5169%	77.8223%	13.1275%	80.5793%	0.0282	0.0398	0.2097
Eddy-Viscosity [69]	7.6400%	48.3684%	8.7583%	56.5878%	0.0276	0.0282	0.1462
Single-state [30]	12.5323%	78.6410%	13.1052%	81.2461%	0.0280	0.0410	0.2111
Our Method	7.4776%	20.4176%	7.8706%	22.7046%	0.0284	0.0272	0.0849

Table 5: **Error on Different Statistics: NS equation** Header: From left to right: Average relative error on energy spectrum, max relative error on energy spectrum, total variation distance from (ground truth) vorticity distribution, average component-wise TV distance(error), max component-wise TV distance(error), and variance of vorticity.

Method	Avg. Eng.	Max Eng.	Vorticity	Avg. TV	Max TV	Variance
CGS (No closure)	178.4651%	404.9923%	0.1512	0.4914	0.8367	253.4234%
Smagorinsky [14]	52.9511%	120.0723%	0.0483	0.2423	0.9195	20.1740%
Single-state [30]	205.3709%	487.3957%	0.1648	0.5137	0.8490	298.2027%
Our Method	5.3276%	8.9188%	0.0091	0.0726	0.2572	2.8666%

H.2 Experiment Results

The error of all statistics we considered for KS equation is listed in Table 4 and plotted in Figure 3.

The error of all statistics we considered for NS equation is listed in Table 5 and plotted in Figure 6.

The visualization of TV error for each (marginal) distribution is shown in Figure 4 and its log scale visualization is shown in Figure 5.

H.3 Ablation Study

We carry out an ablation study for KS equation to verify the effect of perturbing with data loss and CGS data loss. The training dataset is described in Appendix F. The results are as in Figure 7.

We conclude that pretraining with data loss is beneficial to the optimization of the PDE loss function and that pretraining with CGS data can improve the generalization property of the model. Even though CGS data is potentially incorrect, it is utilizable for training because they contain some information of the underlying PDE.

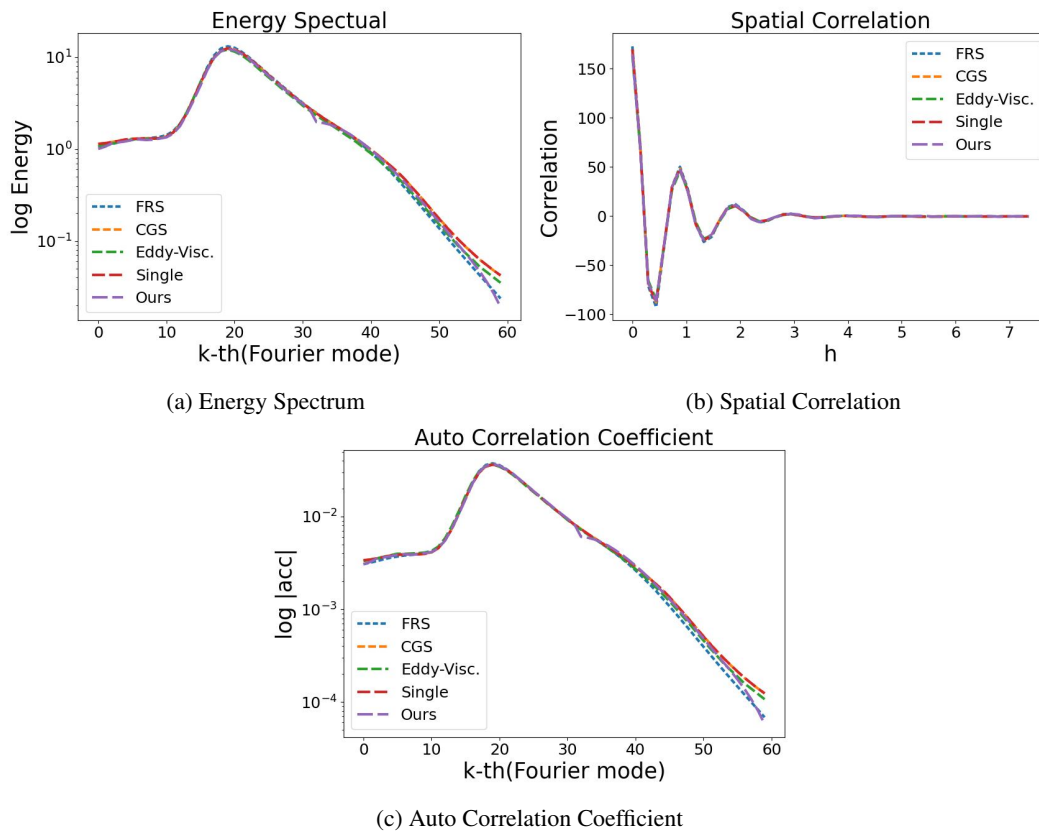


Figure 3: **Experiment Results for KS Equation** 'FRS' (blue line) refers to fully-resolved simulation, and serves as ground truth. 'CGS': coarse-grid simulation(no closure model). 'Eddy-Visc.': classical eddy-viscosity model. 'Single': learning-based single-state closure model. Our method (purple) is closest to ground truth among all coarse-grid methods.

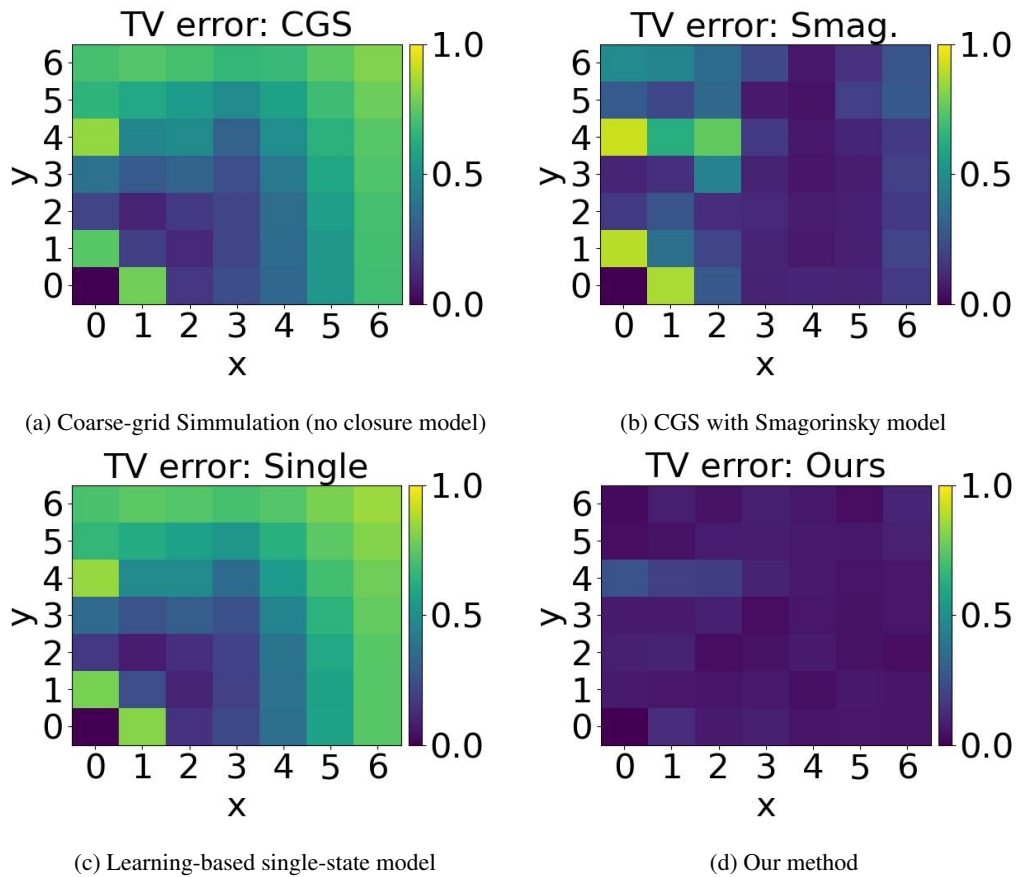


Figure 4: **TV error for NS Equation** The (k, j) -element represents the TV error regarding the distribution of the mode length of (k, j) Fourier basis $e^{i\frac{2\pi}{L}(kx+jy)}$.

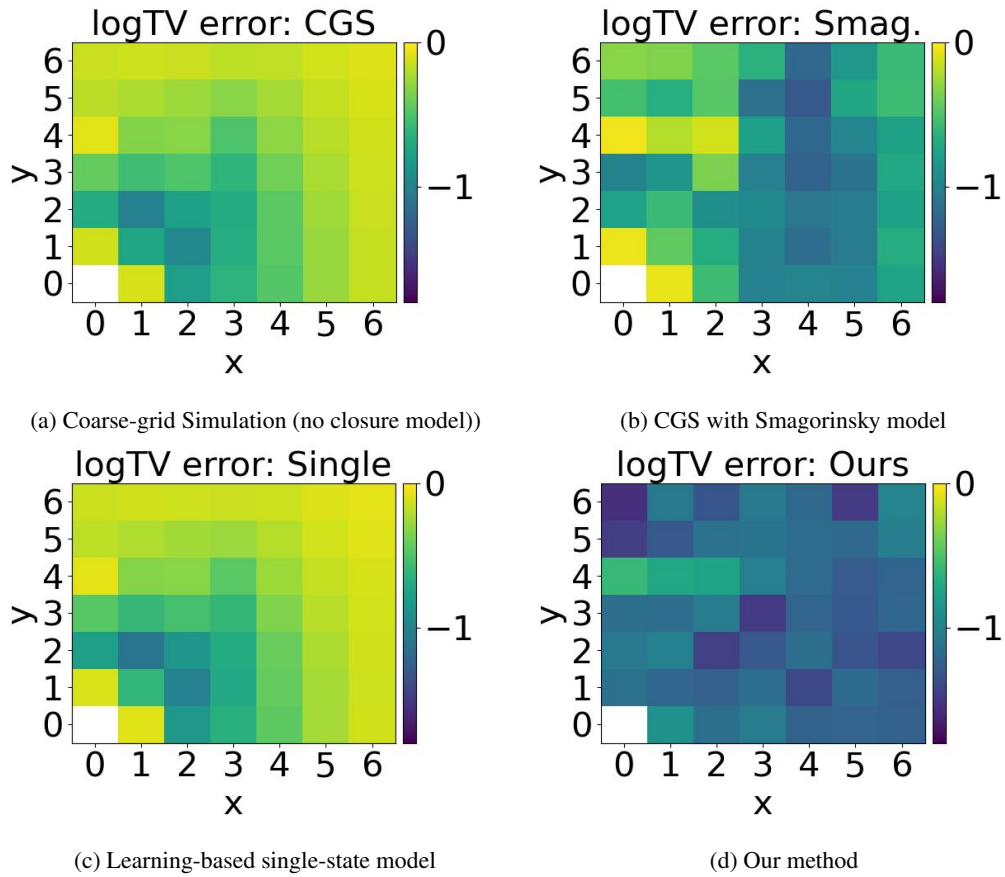


Figure 5: **log-scale TV error for NS Equation** The (k, j) -element represents the logarithm of TV error regarding the distribution of the mode length of (k, j) Fourier basis $e^{i\frac{2\pi}{L}(kx+jy)}$.

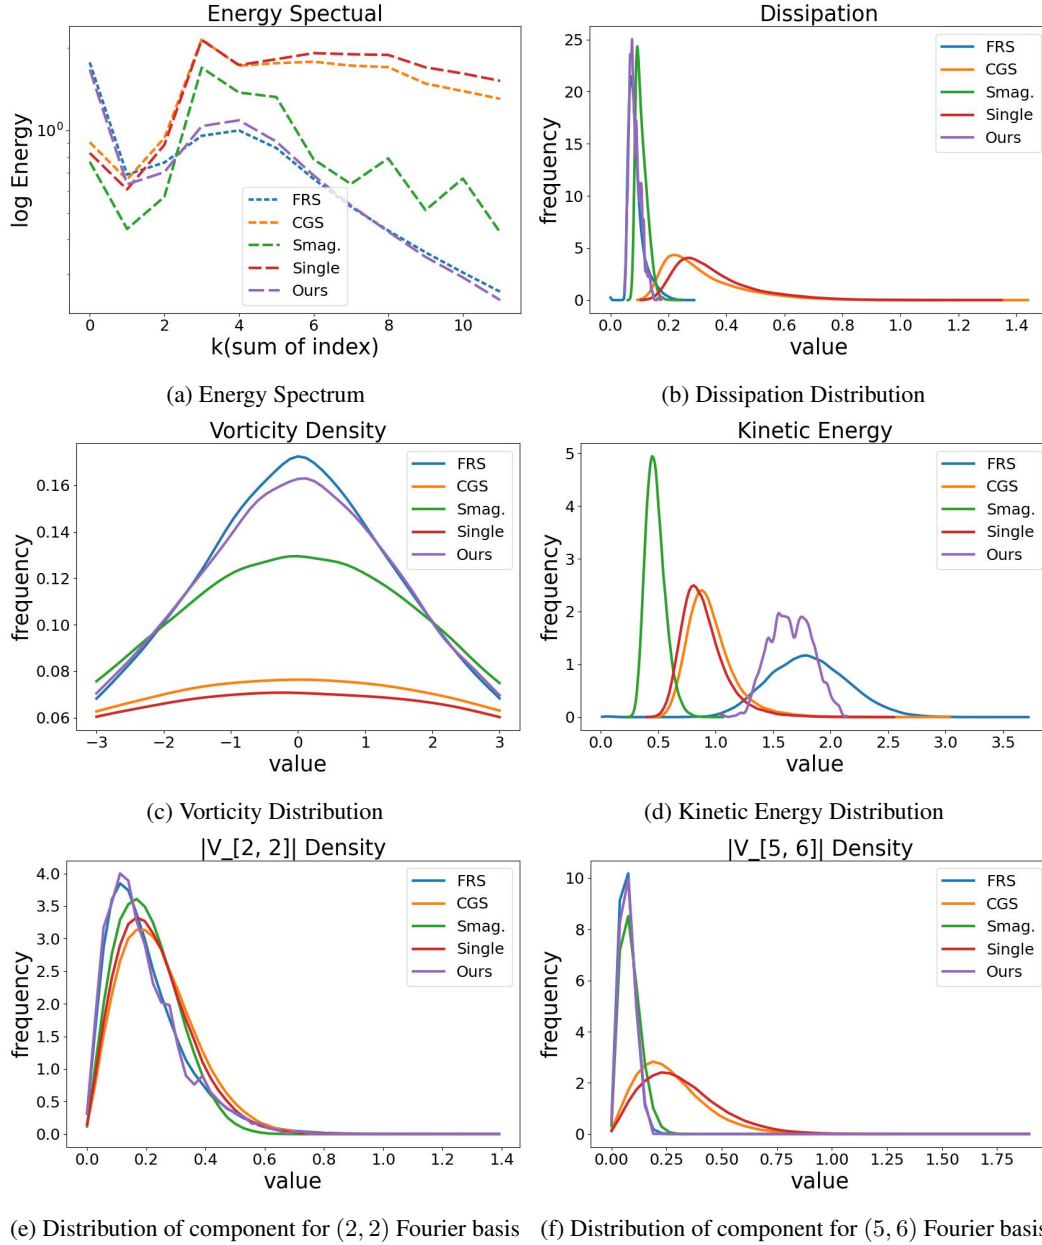


Figure 6: **Experiment Results for NS Equation** 'FRS' (blue line) refers to fully-resolved simulation, and serves as ground truth. 'CGS': coarse-grid simulation(no closure model). 'Smag.': classical Smagorinsky model. 'Single': learning-based single-state closure model. Our method (purple) is closest to ground truth among all coarse-grid methods.

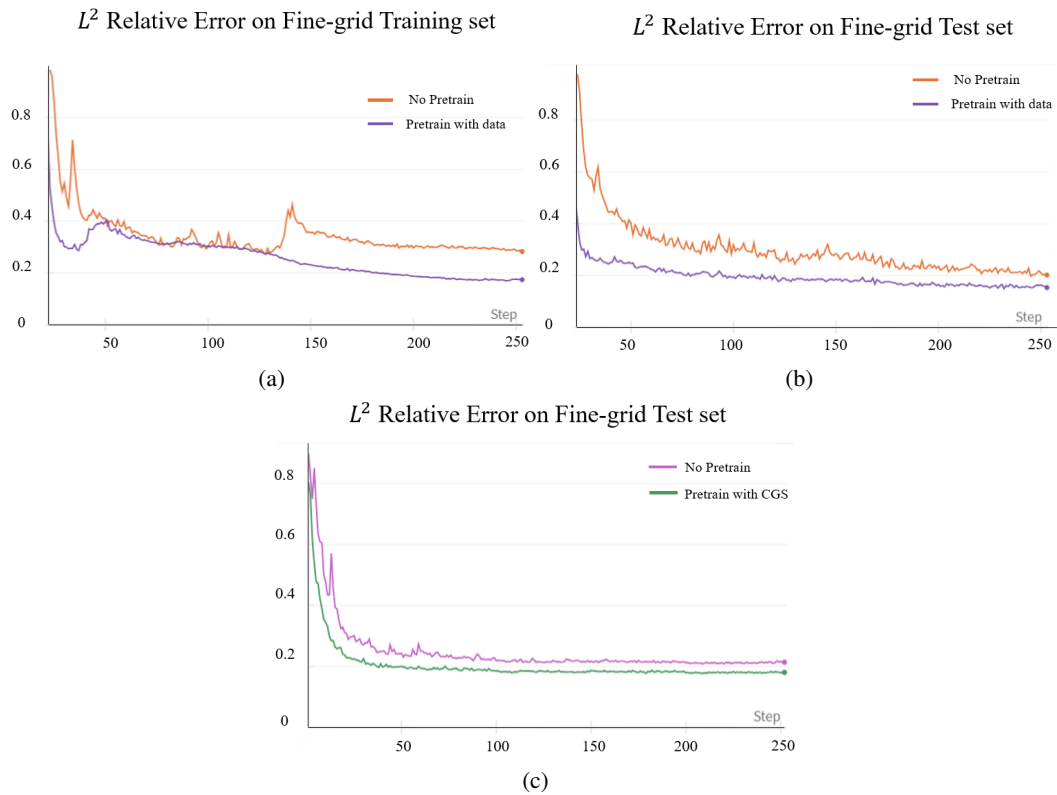


Figure 7: **Experiment Results during training for KS Equation** (a) L^2 relative error on training set for PDE-loss minimization. The optimization achieves a smaller loss when the model has been pre-trained with data loss (purple curve). (b) L^2 relative error on the test set for PDE-loss minimization. The model achieves a smaller error when it has been pre-trained with data loss (purple curve). (c) L^2 relative error on the test set for FRS data-loss minimization. The model has better generalization if it has been pre-trained with CGS data loss (green curve).

D3S3@NeurIPS Paper Checklist (Optional)

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and precede the (optional) supplemental material. The checklist does NOT count toward the page limit.

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Section 1, Appendix A.4

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [NA] .

Justification: The theoretical analysis and method in this manuscript work for general settings (i.e., general dynamics as in Eqn(1)). As for all deep learning methods, we cannot yet provide rigorous guarantees on the convergence of training. Moreover, certain extensions have been out of scope for the current paper, e.g., large-scale 3D equations. These two aspects might be considered our limitations, which we will mention in our final version.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.

- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Appendix B, Appendix D.2, Appendix C, Appendix D, Appendix E.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Appendix G, Appendix H

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.

- (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: The file size of code and dataset exceeds the limit. We will make our code and dataset publicly available after peer-review.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Appendix H

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: Error bars are not reported because it would be too computationally expensive, and they will be hardly noticeable.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Appendix G, Appendix H.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: References

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.

- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.