

SPECTRAL-REFINER: ACCURATE FINE-TUNING OF SPATIOTEMPORAL FOURIER NEURAL OPERATOR FOR TURBULENT FLOWS

Anonymous authors

Paper under double-blind review

ABSTRACT

Recent advancements in operator-type neural networks have shown promising results in approximating the solutions of spatiotemporal Partial Differential Equations (PDEs). However, these neural networks often entail considerable training expenses, and may not always achieve the desired accuracy required in many scientific and engineering disciplines. In this paper, we propose a new learning framework to address these issues. A new spatiotemporal adaptation is proposed to generalize any Fourier Neural Operator (FNO) variant to learn maps between Bochner spaces, which can perform an arbitrary-lengthed temporal super-resolution for the first time. To better exploit this capacity, a new paradigm is proposed to refine the commonly adopted end-to-end neural operator training and evaluations with the help from the wisdom from traditional numerical PDE theory and techniques. Specifically, in the learning problems for the turbulent flow modeling by the Navier-Stokes Equations (NSE), the proposed paradigm trains an FNO only for a few epochs. Then, only the newly proposed spatiotemporal spectral convolution layer is fine-tuned without the frequency truncation. The fine-tuning loss function uses a negative Sobolev norm for the first time in operator learning, defined through a reliable functional-type a posteriori error estimator whose evaluation is exact thanks to the Parseval identity. Moreover, unlike the difficult nonconvex optimization problems in the end-to-end training, this fine-tuning loss is convex. Numerical experiments on commonly used NSE benchmarks demonstrate significant improvements in both computational efficiency and accuracy, compared to end-to-end evaluation and traditional numerical PDE solvers **under certain conditions**.

1 INTRODUCTION

Recently, Deep Learning (DL) pipelines have proven particularly effective in addressing problems formulated by Partial Differential Equations (PDEs). In this paper, we study the problem of learning neural operators (NOs) between infinite-dimensional function spaces (Boullé & Townsend, 2023; Kovachki et al., 2023; Azizzadenesheli et al., 2024; Hoop et al., 2022; de Hoop et al., 2023). Specifically, the problem in consideration is for the Navier-Stokes Equations (NSE) in the turbulent regime ($Re = \mathcal{O}(10^3)$ to $\mathcal{O}(10^4)$). In computation, the difficulties of solving NSE in this regime roots from its “stiffness” attributed to the nonlinear convection with a nearly singular viscous diffusion, as well as the spatiotemporal nature of being highly transient. For this problem, we propose to synergize operator learning with classical numerical PDE methods, complementing one’s drawbacks and limitations with the other’s strengths.

Compromises and drawbacks of traditional numerical methods. To solve NSE, traditional time stepping schemes include as Adams-Bashforth/-Moulton or Runge-Kutta (RK) families (e.g., see Canuto et al. (2007, Appendix D), Karniadakis et al. (1991)). If one opts to use an explicit scheme, or there exists a certain portion of the forcing terms (e.g., the convection term in NSE) computed via explicit schemes, then extremely small time step sizes ($\Delta t \sim \mathcal{O}((\Delta x)^\alpha)$, $\alpha \geq 1$, see e.g., Rannacher (2000, Chapter 4)) must be used, and its necessity is often referred to as the “stiffness” of the PDE. The threshold or constraint of the time step is called the Courant-Friedrichs-Lewy (CFL) condition (e.g., see Johnston & Liu (2004); Wang (2012); He & Sun (2007)). CFL poses a sufficient condition on the step size for “stability”, and this requires the step size usually much smaller than

the requirement for “consistency”. Here, this temporal consistency usually refers to a first-order optimal local truncation error, e.g., how the original Butcher tableau is derived for RK (Butcher, 1965). The CFL puts a threshold for *all* explicit time-marching schemes on how fast the local errors in different frequencies can propagate, and thus prevent the accumulation “pollutes” the approximation globally to ensure stability. This means that, for any time-marching schemes, finer mesh (better spatial consistency) requires the time steps to be much smaller to prevent errors from traveling to neighboring nodes and elements, which greatly increases the computational cost.

Spatiotemporal operator learning. Among the end-to-end operator learning for NSE, the common approach is the so-called autoregressive “roll-out”. During rolling-out, several snapshots of solutions are concatenated as the input to the NO, which outputs an approximated solution at the subsequent time step. This procedure can be repeated recurrently until reaching the model’s stability capacity. **In contrast to the traditional time marching solvers whose step sizes are restricted by the CFL condition, roll-out can withstand a much bigger time step size.** NOs used in roll-out approach include the original FNO2d (Li et al., 2020), and those in Li et al. (2022); Brandstetter et al. (2023); Gupta & Brandstetter (2023); Fonseca et al. (2023). However, for an end-to-end operator learner, the roll-out approach faces the same dilemma of super-resolution capacity in the temporal dimension as a traditional solver: finer time steps cost prohibitively more, while larger time steps does not guarantee stability. To solve this problem, we turn to the prevalent functional analytic framework for studying the convergence and stability of solution trajectories for NSE: the theory of Bochner spaces, e.g., Aubin-Lions lemma (Lions & Magenes, 2012, Chapter 3), (Temam, 1995, Chapter 2), (Evans, 2022, Chapter 7). Inspired by this theoretical insight, we propose a Spatiotemporal adaption for all Fourier Neural Operator (ST-FNO) variants, which now can perform arbitrary zero-shot super-resolution not only spatial dimensions, but also allows the temporal dimension to vary for the first time. For a prototype spatiotemporal operator learner FNO3d in Li et al. (2020), it learns a map between a fixed number of spatial “snapshots” (product spaces). The newly proposed ST-FNO directly learns a map between Bochner spaces $L^2(\mathcal{T}_0; \mathcal{V})$ to $L^2(\mathcal{T}; \mathcal{V})$ on non-overlapping time intervals \mathcal{T}_0 and \mathcal{T} . This makes the model become “trajectory-to-trajectory”, where the initial trajectory is the input of an NO to obtain the output evaluation as an approximation to the subsequent trajectory of the solution. Here \mathcal{V} denotes the spatial Hilbert space in which snapshots of the solution at a specific time reside, and for a detailed notation list, we refer the readers to Appendix A.

Limits and the lack of accuracy for neural operators. The NO approach has the potential to bypass various difficulties and compromises of traditional schemes of numerical PDEs. However, in all end-to-end operator learning benchmarks of NSE, even the state-of-the-art NOs still fall short in achieving high-accuracy solutions. For example, an end-to-end roll-out approach suffers from error accumulation experimentally (e.g., see Li et al. (2022, Figure 9)). To our best knowledge, no NO-only-based operator learning approach can break the barrier of a 2-digit accuracy in terms of the relative difference with the ground truth under the Bochner norm. Moreover, to our best knowledge, NO-only approaches have no theoretical stability estimate, such as the error propagation operator is a contraction. Recently, a remarkable advancement called PDE-refiner (Lippe et al., 2023) learns an extra error correction NO under the Denoising Diffusion framework (Ho et al., 2020). For a single instance, it can get a stable long roll-out, and achieve for the first time $\mathcal{O}(10^{-8})$ relative difference with the ground truth after a single time marching step, and $\mathcal{O}(10^{-6})$ in the Bochner norm. Nevertheless, for all models above and their learning frameworks, the optimization is to minimize the difference between the outputs from the NO, namely $u_{\mathcal{N}}$, and the ground truth $u_{\mathcal{S}}$, generated by a traditional numerical PDE solver. The difference with the true solution u under a certain norm is not directly optimized, as the analytical expression of u is unknown most of the time in real-life applications. For difficult PDEs such as the NSE, in a linear time-stepping scheme, the ground truth $u_{\mathcal{S}}$ may on its own only have 3-digit accuracy in the Bochner norm. Note that this may already be of optimal order $\mathcal{O}(\Delta t)$ by convergence estimates (Heywood & Rannacher, 1986). Therefore, minimizing the difference between $u_{\mathcal{N}}$ and $u_{\mathcal{S}}$ becomes fruitless if the numerical approximation (ground truth) $u_{\mathcal{S}}$ is not accurate at the first place, as unnecessary computational resources may have been spent to get closer to $u_{\mathcal{S}}$.

New hybrid scheme. To address these dilemmas for NSE’s approximation, we take an alternative hybrid learning paradigm inspired by traditional numerical methods. Unlike the arduous training of running hundreds of epochs, the newly proposed ST-FNO is trained for only a few epochs (e.g., 10), concluding with the freezing of most model parameters. Then, the *last* new spatiotemporal spectral convolution layer of ST-FNO, which is attributed to its temporal arbitrary-lengthed inference capacity, is fine-tuned with the help of traditional solvers. During fine-tuning, this layer is relieved from the

common frequency truncation FNO has. A traditional solver with a single time step is used to obtain highly accurate approximations for extra field variables to conform with the physics. These extra fields help the evaluation of a new loss function in fine-tuning, a functional-type *a posteriori* error estimate measured under a *negative* Sobolev norm, which is equivalent to the variational form-associated norm for the NSE. Note that there shall be no training of extra error-correcting NOs (Dresdner et al., 2023; Lippe et al., 2023). Moreover, unlike the same nonconvex optimization problem that the extra error-correcting NOs try to tackle, our fine-tuning optimization problem is *convex*, and allows us to achieve high accuracy with a fractional of computational resources when compared with other fine-tuning approaches that optimize the whole model with a physics-informed loss.

Fine-tuning using functional-type *a posteriori* error estimation. To seek highly accurate NO approximated solutions that are closer to the true solution directly, we turn to the *a posteriori* error estimation techniques. This technique allows computing the error without knowing the true PDE solutions, and has been widely studied for Galerkin-type methods, such as finite element (Ainsworth & Oden, 1993; 1997; Oden et al., 1994), as well as for parameterized PDEs (Hesthaven et al., 2016; Patera et al., 2007). Among all types of *a posteriori* error estimation techniques, functional-type *a posteriori* estimator (Repin, 2008) views the error as a functional on the test Sobolev spaces and evaluates accurate representations with the help of an extra dual variable (Ern & Vohralík, 2010). In our hybrid approach, inspired by this, we combine the strengths of NOs and traditional numerical solvers. Using the newly proposed spatiotemporal discretization-invariant NO, we propose to use a negative Sobolev norm in the frequency domain as a functional-type *a posteriori* error estimation as the fine-tuning loss. Traditionally, the *a posteriori* error estimation is used to refine local basis, yet this constraint on the purpose of the method leads to inaccurate localized representations for the H^{-1} error functional. Our new approach is not attached to the local refining requirement Bonito et al. (2024). As a result, the negative Sobolev norm used in the new loss is handily defined through the Gelfand triple in the frequency domain which is global. Our method needs no extra dual variable reconstruction problem as in the traditional methods (Ern & Vohralík, 2010). Unlike commonly adopted physics-informed operator learning (Li et al., 2024c), NO prediction for other field variables are not needed either. The extra field variables for computing the error are recovered through an auto-differentiable numerical solver with the NO output of the primal variable (vorticity or velocity). Overall, this practice leads to “refining” the learnable set of the global spectral basis.

For a more detailed review of operator learning and further motivation discussion with a much higher degree of mathematical rigor, we refer the readers to Appendix B.

Main contributions. The main contributions of this work are summarized as follows. For the difference between the common roll-out approach versus direct spatiotemporal learning between Bochner spaces, we refer the readers to Figure 1.

- **Spatiotemporal Fourier Neural Operator.** We design the first spatiotemporal adaption technique for all FNO variants (ST-FNO) to enable them to learn maps between Bochner spaces.
- **New hybrid operator learning paradigm.** We propose to train and evaluate ST-FNOs using a new strategy for spatiotemporal predictions, which has significantly improved over the existing methods in speed and accuracy. Only a few epochs of training combined with fine-tuning yield highly accurate approximations, extra field variables are obtained by auto-differentiable multistep solvers (such as RK4), the input of which are neural predictions. The auto-differentiability of the solver makes it possible for the fine-tuning to optimize the parameters in FNO.
- **Functional *a posteriori* error estimation.** Leveraging the spectral structure of the new ST layer in ST-FNO, Parseval identity, and Gelfand’s triple, the fine-tuning utilizes a new loss, *a posteriori* error estimation in the negative Sobolev norm (functional norm). This procedure does not require any ground truth data (e.g., numerical solution of the underlying PDEs generated by traditional numerical solvers), nor the knowledge of the true solution(s) to the PDE. This new loss is proven to be reliable in theory, in the meantime much more efficient in the ablation experiments.
- **Experimental results.** We created a native PyTorch port of Google’s Computational Fluid Dynamics written in JAX (Kochkov et al., 2021; Dresdner et al., 2023) with enhanced functionality for tensor operations such as the facilitation of fine-tuning for the latent fields, publicly available at <https://anonymous.4open.science/r/torch-cfd-E6B0>, with scripts to replicate the experiments as well as the data generation. The data are available at <https://www.kaggle.com/datasets/anonymousauthor25/sfno-dataset>.

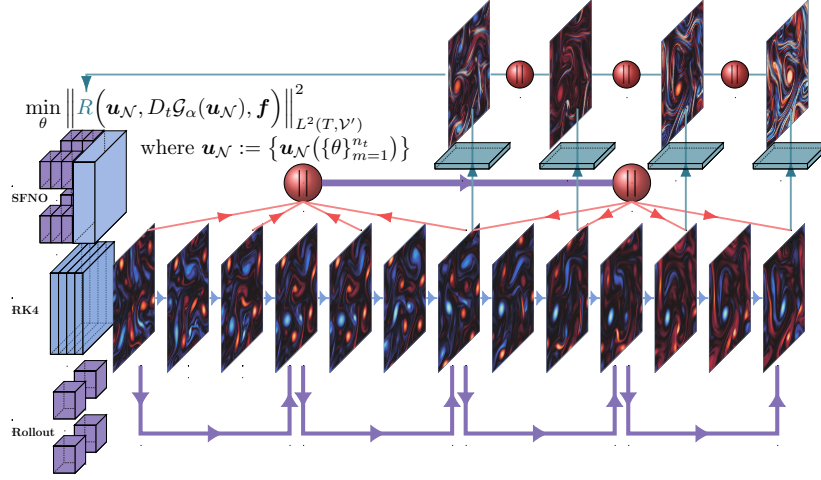


Figure 1: Schematic differences between approaches. 4th-order Runge-Kutta (RK4): small time steps bounded by the CFL condition, de-aliasing filter needed; autoregressive NO rolling-out: using previous evaluation as input repetitively, large time steps, no stability guarantees; Spatiotemporal FNO (ST-FNO) with hybrid fine-tuning: large time steps, yielding arbitrary-lengthed temporal prediction in a single evaluation, fine-tuning stability inheriting from traditional solvers, no de-aliasing filter preserves all higher modes, parallel-in-time optimization.

2 SPATIOTEMPORAL OPERATOR LEARNING FOR NAVIER-STOKES EQUATIONS

Both drawbacks and advantages of traditional numerical methods and NO-based methods in Section 1 play vital roles in shaping our study. We first briefly discuss the spatiotemporal operator learning problem on Bochner spaces associated with NSE. Then, we detail how to modify a generic FNO architecture to obtain an operator learner between Bochner spaces.

2.1 SPATIOTEMPORAL OPERATOR LEARNING PROBLEM FOR NSE

For 2D NSE, the velocity field $\mathbf{u}(t, \mathbf{x}) : [0, T] \times \Omega \rightarrow \mathbb{R}^2$ is seen as an element in the Bochner space $L^p([0, T], \mathcal{V})$ where \mathcal{V} is a spatial Sobolev space in which each snapshot at t of the solution $\mathbf{u}(t, \cdot)$ resides. As in Li et al. (2020), we consider the vorticity-streamfunction (V-S) formulation (2.1) with periodic boundary condition (PBC). We also consider the velocity-pressure (V-P) formulation (2.2). In $\Omega \times (0, T]$, for vorticity $\omega := \nabla \times \mathbf{u}$, and streamfunction ψ , these two formulations read

$$\text{(Vorticity-Streamfunction)} \quad \partial_t \omega + \mathbf{rot} \psi \cdot \nabla \omega - \nu \Delta \omega = \nabla \times \mathbf{f}, \quad \omega + \Delta \psi = 0. \quad (2.1)$$

$$\text{(Velocity-Pressure)} \quad \partial_t \mathbf{u} + (\mathbf{u} \cdot \nabla) \mathbf{u} - \nu \Delta \mathbf{u} + \nabla p = \mathbf{f}, \quad \nabla \cdot \mathbf{u} = 0. \quad (2.2)$$

For all analyses in line with the Hilbertian framework, $\mathcal{V} = H^1(\mathbb{T}^2)$ for vorticity and $\mathbf{H}^1(\mathbb{T}^2)$ for velocity, where \mathbb{T}^2 denotes the unit torus, i.e., $\Omega = (0, 1)^2$ with a component-wise PBC. For a fixed forcing function in \mathcal{V}' , the initial condition is either $\omega(0, \mathbf{x}) = \omega_0(\mathbf{x})$ or $\mathbf{u}(0, \cdot) = \mathbf{rot}((-\Delta)^{-1} \omega_0)$. Here ω_0 is drawn from a compactly supported probability measure μ , in which the compactness corresponds to certain power/energy spectrum decay law to produce isotropic turbulence (McWilliams, 1984), and we refer the reader to Appendix C for details. Then, we can consider the following map $G^\mu : \prod_{i=1}^\ell \mathcal{V} \rightarrow \prod_{i=1}^{n_t} \mathcal{V}$:

$$G^\mu : \mathbf{a} := [\omega(t_1, \cdot), \dots, \omega(t_\ell, \cdot)] \mapsto \mathbf{u} := [\omega(t_{\ell+1}, \cdot), \dots, \omega(t_{\ell+n_t}, \cdot)], \quad (2.3)$$

where $t_{\ell+n_t} \leq T$, and the input and the output are snapshots obtained from a solution trajectory with the same initial condition ω_0 . The operator learning task for NSE using a prototype spatiotemporal operator learner, e.g., FNO3d (Li et al., 2020), is to learn this G^μ between a fixed number of Cartesian products of spatial Sobolev spaces. Specifically, $G^\mu : \mathcal{X} \rightarrow \mathcal{Y}$ maps elements in $\mathcal{X} = \prod_{i=1}^\ell \mathcal{V}$ to elements in $\mathcal{Y} = \prod_{i=1}^{n_t} \mathcal{V}$. In the task for this prototype “snapshots learner”, the number of snapshots $\ell \in \mathbb{Z}^+$ is **fixed**. As such, \mathcal{X} and \mathcal{Y} represent spaces of two non-overlapping temporal segments of

solution snapshots. Based on these snapshots’ discretized approximations from data, these snapshots learners learn an operator $G_\theta : \mathcal{X} \rightarrow \mathcal{Y}$, where the number of parameters in G_θ is independent of the spatial discretization size, yet does depend on the number of snapshots ℓ . This dependence in the setting of this task makes it not “trajectory-to-trajectory”.

In this study, the learning aims to recast (2.3) to a trajectory-to-trajectory operator learning problem, conforming to the Hilbertian formulation of NSE using Bochner spaces. Using the fact that the weak solutions to (2.1) and (2.2) at a given time interval \mathcal{T} is in $L^2(\mathcal{T}; \mathcal{V})$, the operator to be learned is

$$\tilde{G} : L^2(t_1, t_\ell; \mathcal{V}) \rightarrow L^2(t_{\ell+1}, t_{\ell+n_t}; \mathcal{V}) \text{ for } \mathcal{V} := H^1(\mathbb{T}^2) \text{ or } \{\mathbf{v} \in \mathbf{H}^1(\mathbb{T}^2) : \nabla \cdot \mathbf{v} = 0\}. \quad (2.4)$$

In what follows, we shall present how to modify any common FNO variant to become a Bochner space operator learner that can learn maps from arbitrary-sized discretization in $\mathbb{R}^{\ell \times n \times n \times d}$ to $\mathbb{R}^{n_t \times n \times n \times d}$ ($d = 1$ in V-S; $d = 2$ in V-P). This operator can be trained by a standard end-to-end supervised learning pipeline using lower-resolution data, and for inference, the newly proposed spatiotemporal trajectory-to-trajectory learner, ℓ, n_t, n can all be of variable sizes. In contrast, the snapshot learner FNO3d predicts the subsequent n_t snapshots with the first ℓ snapshots as input, with both n_t and ℓ fixed, and only the spatial resolution n can be much higher than the input-output pairs used in the training. Hereafter we omit the d dimension if no ambiguity arises.

2.2 SPATIOTEMPORAL ADAPTATION OF FOURIER NEURAL OPERATORS

For any FNO variant, such as FNO3d (Li et al., 2020) or Factorized FNO (Tran et al., 2023), with the following meta-architecture: $G_\theta := Q \circ \sigma_L \circ K_{L-1} \circ \dots \circ \sigma_1 \circ K_0 \circ P$, where P is a lifting operator, Q is a projection operator that does a pointwise reduction in the channel dimension, σ_j can be a pointwise universal approximator or simply chosen as a nonlinearity. $K_j := K_{\phi_j}$ is the spectral convolution that does a spatiotemporal 3D FFT. During training, the operator to be learned is restricted to finite-dimensional subspaces $\mathcal{X} \supset \mathcal{X}_n \rightarrow \mathcal{Y}_n \subset \mathcal{X}$, in the sense that the continuous spatial Sobolev space \mathcal{V} in the product spaces is replaced by a finite-dimensional subspace $\mathcal{V} \supset \mathcal{S} \simeq \mathbb{R}^{n \times n}$ with continuous embeddings $\{\mathbf{a}_S \in \mathbb{R}^{\ell \times n \times n}\} \hookrightarrow \mathcal{X}$, and $\{\mathbf{u}_S \in \mathbb{R}^{n_t \times n \times n}\} \hookrightarrow \mathcal{Y}$. The positional encodings $\mathbf{p} := (t_i, \mathbf{x}_j)_{i=1}^\ell \in \mathbb{R}^{3 \times \ell \times n \times n}$ represents the spatiotemporal grid, and is concatenated to \mathbf{a}_S before feeding to P .

For spatiotemporal problems, all FNO variants novelly exploit a convenient architectural advantage of operator-valued NNs: the input temporal dimension ℓ is treated as the channel dimension of an image, thus enabling channel mixing as a learnable temporal extrapolation. However, this neat trick coincidentally makes the lifting operator the biggest hurdle for G_θ to become a trajectory-to-trajectory operator learner between Bochner spaces, since P ’s dimension must be hard-coded¹ and thus depends on the input pair’s time steps.

In what follows, we use FNO3d as an example to present three new modifications to FNO3d to become a trajectory-to-trajectory learner ST-FNO3d, that is, to act as an operator that maps an arbitrary-time-step input to an arbitrary-time-step output. We also note that this modification applies to any FNO variant that predicts spatial and temporal dimensions at the same time. These changes are so simple that the trajectory-to-trajectory modification can serve as drop-in replacements for their snapshot learner counterparts when the temporal input dimensions are fixed. For the schematic difference, please refer to Figure 2 and 3. For more technical details, we refer the readers to Appendix D.2.

Modifications to 3D FFT. The first change is in the spatiotemporal FFT from an analytical point of view. Using the spectral convolution $K(\cdot)$ in FNO3d as an example, with a slight abuse of notation, denote the latent dimension (width) by d_v , let the vector-valued latent representation with d_v channels be continuously embedded into $\prod_{i=1}^{n_t} \mathcal{V}$ by $\mathbf{v}_S \hookrightarrow \mathbf{v}$ in each channel, and denote $\Lambda := [t_{\ell+1}, \dots, t_{\ell+n_t}]$, then $K(\cdot)$ is “semi-discrete” in a sense as follows

$$(K\mathbf{v})(t, \mathbf{x}) := (W\mathbf{v})(t, \mathbf{x}) + \sum_{s \in \Lambda} \int_{\Omega} \kappa((t, \mathbf{x}), (s, \mathbf{y})) \mathbf{v}(s, \mathbf{y}) dm(\mathbf{y}), \quad (2.5)$$

where t takes values only in a discrete set Λ , and $\mathbf{x} \in \Omega$. Here $W \in \mathbb{R}^{(d_v+1) \times d_v}$ is a pointwise-applied affine linear operator, $\kappa \in C((\Lambda \times \Omega) \times (\Lambda \times \Omega); \mathbb{R}^{d_v \times d_v})$, and m denotes an approximation

¹Line 99 in `fourier_3d.py` in the master branch of `gh/neuraloperator/neuraloperator`




Figure 2: The FNO3d in Li et al. (2020) is a snapshot learner. ■: spectral convolution layer ; ■: pointwise nn.Conv3d that works as channel expansion/reduction; ■: pointwise nonlinearity. The TikZ source code to produce this figure is modified from the examples in Iqbal (2018).




Figure 3: The Spatiotemporal-adaptated FNO3d (ST-FNO3d) is now a trajectory-to-trajectory learner. ■: layer normalization to replace a hard-coded global normalization. Combined with channel mixing, the first spectral convolution layer serves as a time-depth-wise separable (global) convolution, after which the time dimension is shrunk to a fixed “latent” time dimension through iFFT’s resampling. ■: the spatiotemporal spectral convolution layer as the final layer is fine-tuned after the training phase.

to the Lebesgue measure on Ω . Evolving K into a spatiotemporal spectral convolution acting on Bochner space is straightforward. We change (2.5) slightly by adopting of an atom-like measure $\delta(\cdot)$ in the temporal dimension, which then generalizes to the variable-length temporal discretization for any $(t, \mathbf{x}) \in (a, b) \times \Omega$, i.e., one can obtain arbitrarily many snapshots on the interval of interest

$$(\tilde{K}(\mathbf{v}))(t, \mathbf{x}) := (W\mathbf{v})(t, \mathbf{x}) + \int_a^b \int_{\Omega} \kappa((t, \mathbf{x}), (s, \mathbf{y})) \mathbf{v}(s, \mathbf{y}) dm(\mathbf{y}) d\delta(s). \quad (2.6)$$

During training using the discretized data, each hidden layer is a map from $\mathbb{R}^{d_v \times d_t \times n \times n} \rightarrow \mathbb{R}^{d_v \times d_t \times n \times n}$, where d_t is a “latent” dimension of time steps that is chosen $\leq n_t$. The layerwise propagation mechanics remains the same: $\tilde{K}_{\phi} \mathbf{v} := W\mathbf{v} + \mathcal{F}^{-1}(R_{\phi} \cdot (\mathcal{F}\mathbf{v}))$, where \mathcal{F} and \mathcal{F}^{-1} denote the spatiotemporal Fourier transform and its inverse, respectively. The global spatiotemporal interaction characterized by the kernel is truncated in terms of modes in the frequency domain at (τ_{\max}, k_{\max}) , such that $\{(\tau, \mathbf{k}) \in \mathbb{Z}^3 : |\tau| \leq \tau_{\max}, |\mathbf{k}_j| \leq k_{\max}, j = 1, 2\}$ are kept. R_{ϕ} is parametrized as a $\mathbb{C}^{\tau_{\max} \times k_{\max} \times k_{\max} \times d_v \times d_v}$ -tensor. Here (τ, \mathbf{k}) denotes the coordinate in the spatiotemporal Fourier domain. Note that the integral represented by matrix product with $\mathcal{F}\kappa(\cdot, (s, \mathbf{y}))$ can then be viewed as residing in the continuous space as the Fourier basis (3.1) can be evaluated at arbitrary point.

Modifications to the lifting operator P . The other major hurdle for FNO3d to become trajectory-to-trajectory is that a global normalization is applied with a hard-coded temporal dimension using training data.² In ST-FNO3d’s modification, a periodic padding along the temporal dimension is applied to the input, and then a time-depth separable convolution layer \mathcal{I} with variable time steps is used to map an arbitrary number of snapshots to a fixed number of channels d_v with a fixed latent time steps d_t . The spatiotemporal positional encodings $\tilde{\mathbf{p}}_S := W_p \mathbf{p}$, with W_p a learnable random projection with periodically padded \mathbf{p} as input. This makes $\tilde{\mathbf{p}}_S$ match the latent fields’ dimensions such that they can be concatenated. The latent fields, concatenated with $\tilde{\mathbf{p}}_S$, are then normalized by a learnable layer normalization layer $\text{Ln}(\cdot)$, instead of a global fixed pointwise normalizer for the raw data in FNO3d. Finally, in ST-FNO3d, the lifting operator becomes $\tilde{P} := \text{Ln}(\tilde{\mathbf{p}}_S + \mathcal{I}(\cdot)) : \mathbb{R}^{\ell \times n \times n} \rightarrow \mathbb{R}^{d_v \times d_t \times n \times n}$.

²Line 220 and 224 in `fourier_3d.py` in the master branch of [gh/neuraloperator/neuraloperator](https://github.com/neuraloperator/neuraloperator)

Modifications to the projection operator Q . In FNO3d, the original out projection operator $Q : \mathbb{R}^{n_t \times d_v \times n \times n} \rightarrow \mathbb{R}^{n_t \times n \times n}$. In spatiotemporal-adapted $\tilde{Q} : \mathbb{R}^{d_v \times d_t \times n \times n} \rightarrow \mathbb{R}^{n_t \times n \times n}$, which maps the last latent representation to match the dimension of the output \mathbf{u}_S . In \tilde{Q} , the key for an arbitrary-lengthed inference is that we compose another spectral convolution K_S after channel reduction. K_S can be thought of as a post-processing layer (also as the only to be fine-tuned). It takes advantage of the FFT/iFFT’s natural super-resolution capacity, especially in the temporal dimension, by zero-padding the latent temporal step dimension (d_t) to given arbitrary output time steps. For the necessity of this padding, we refer the reader to Figure 7 for an illustrative example. For the V-P formulation, to impose the divergence-free condition for, S is implemented as an optional non-learnable Helmholtz decomposition layer in \tilde{Q} .

3 NEW HYBRID PARADIGM FOR SPATIOTEMPORAL OPERATOR LEARNING

Built upon a reasonably accurate approximation, the fine-tuning of ST-FNO is proposed. It is able to yield accuracy on par with traditional numerical methods on the same time horizon, and computational resources used are comparable to the evaluation of NOs. Taking advantage of the efficient ST-FNO zero-shot arbitrary-lengthed temporal inference, this new approach does not need thousands of marching steps like traditional methods. Meanwhile, it borrows the wisdom from traditional Galerkin methods to improve the accuracy (consistency) of the NO approach, liberating the scheme from trade-offs that the traditional methods must make to ensure stability and convergence.

3.1 A POSTERIORI ERROR ESTIMATION USING A FUNCTIONAL NORM

We shall present the proposed fine-tuning using the V-P formulation in subsequent subsections. Denote an ST-FNO evaluation at a specific $t_m \in \mathcal{T}_{n_t} := \{t_{\ell+1}, \dots, t_{\ell+n_t}\}$ in the output time interval as $\mathbf{u}_{\mathcal{N}}^{(m)}$. The construction of \tilde{Q} in ST-FNO renders $\mathbf{u}_{\mathcal{N}}^{(m)} \in \mathcal{W}$, where \mathcal{W} is the divergence-free subspace of $\mathcal{S}|_{t=t_m} \times \mathcal{S}|_{t=t_m} \subset \mathcal{V} := \mathbf{H}^1(\mathbb{T}^2)$, where

$$\mathcal{S} := \text{span} \{ \Re(e^{i\tau_m t} e^{i\mathbf{k} \cdot \mathbf{x}}) : -n/2 \leq k_j \leq n/2 - 1, -n_t/2 \leq m \leq n_t/2 - 1 \} / \mathbb{R}, \quad (3.1)$$

for $\mathbf{k} = 2\pi(k_j)_{j=1,2}$ and $\tau_m = 2\pi m/(T - t_\ell)$. Then, a temporal continuous approximation $\mathbf{u}_{\mathcal{N}} := \mathbf{u}_{\mathcal{N}}(t, \cdot)$ can be naturally defined by allowing t vary continuously on $\mathcal{T} := [t_{\ell+1} - t_p, T + t_p]$ thanks to the spectral basis of \mathcal{S} , where t_p is the temporal periodic padding in Section 2.2. Define residual functional $R(\mathbf{u}_{\mathcal{N}}) \in L^2(\mathcal{T}; \mathcal{V}')$: at $t \in \mathcal{T}$ and for $\mathbf{v} \in \mathcal{V}$

$$R(\mathbf{u}_{\mathcal{N}}) := \mathbf{f} - \partial_t \mathbf{u}_{\mathcal{N}} - (\mathbf{u}_{\mathcal{N}} \cdot \nabla) \mathbf{u}_{\mathcal{N}} + \nu \Delta \mathbf{u}_{\mathcal{N}}, \text{ and } R(\mathbf{u}_{\mathcal{N}})(\mathbf{v}) := \langle R(\mathbf{u}_{\mathcal{N}}), \mathbf{v} \rangle. \quad (3.2)$$

$R(\mathbf{u}_{\mathcal{N}})$ measures how PDE (2.2) is violated by the finite-dimensional approximation $\mathbf{u}_{\mathcal{N}}$, not in a localized/pointwise fashion, but rather in a global way by representing the error based on its inner product against arbitrary \mathbf{v} . At a specific time t , the functional norm of $R(\mathbf{u}_{\mathcal{N}})(t, \cdot) \in \mathcal{V}'$ defined as follows is then an excellent measure of the error to be a candidate for a loss function in view of Theorem 3.1:

$$\|R(\mathbf{u}_{\mathcal{N}})(t, \cdot)\|_{\mathcal{V}'} := \sup_{\mathbf{v} \in \mathcal{V}, \|\mathbf{v}\|_{\mathcal{V}}=1} |(R(\mathbf{u}_{\mathcal{N}}), \mathbf{v})|. \quad (3.3)$$

Theorem 3.1 (A posteriori error bound for any fine-tuned approximations, informal version). *Let the weak solution to (2.2) be $\mathbf{u} \in L^2(\mathcal{T}; \mathcal{V})$, and $\partial_t \mathbf{u} \in L^2(\mathcal{T}; \mathcal{V}')$. Assume \mathbf{u} be sufficiently regular, then the dual norm of the residual is efficient to estimate the error for any $\mathbf{u}_{\mathcal{N}}$:*

$$\|R(\mathbf{u}_{\mathcal{N}})\|_{L^2(\mathcal{T}; \mathcal{V}')}^2 \lesssim \|\mathbf{u} - \mathbf{u}_{\mathcal{N}}\|_{L^2(\mathcal{T}; \mathcal{V})}^2 + \|\partial_t(\mathbf{u} - \mathbf{u}_{\mathcal{N}})\|_{L^2(\mathcal{T}; \mathcal{V}')}^2. \quad (3.4)$$

Moreover, if \mathbf{u} and $\mathbf{u}_{\mathcal{N}}$ are “sufficiently close”, then it is reliable to serve as an error measure:

$$\|\mathbf{u} - \mathbf{u}_{\mathcal{N}}\|_{L^\infty(\mathcal{T}_m; \mathcal{H})}^2 + \|\mathbf{u} - \mathbf{u}_{\mathcal{N}}\|_{L^2(\mathcal{T}_m; \mathcal{V})}^2 \leq \|(\mathbf{u} - \mathbf{u}_{\mathcal{N}})(t_m, \cdot)\|_{\mathcal{V}}^2 + C \int_{\mathcal{T}_m} \|R(\mathbf{u}_{\mathcal{N}})(t, \cdot)\|_{\mathcal{V}'}^2 dt, \quad (3.5)$$

where $\mathcal{T}_m := (t_m, t_{m+1}]$, and the constants depend on the regularity of the true solution \mathbf{u} .

3.2 FINE-TUNING USING NEGATIVE SOBOLEV NORM AS LOSS

The functional norm (3.3) is “global” because it does not have a natural ℓ^2 -like summation form where each summand can be evaluated in a localized neighborhood of grid points. Nevertheless,

thanks to the Gelfand triple, and viewing the Fourier transform as an automorphism in the tempered distribution space (e.g., see Peetre (1975) and Gel'fand & Shilov (2016, Chapter 3)), we can define the pairing between \mathcal{V} and \mathcal{V}' as follows without getting into the intricate natures of the tempered distribution:

$$\langle f, g \rangle_{\mathcal{V}, \mathcal{V}'} \text{ “=” } \int_{\mathbb{Z}^2} (1 + |\mathbf{k}|^2)^{-s} \hat{f}(\mathbf{k}) \overline{\hat{g}(\mathbf{k})} (1 + |\mathbf{k}|^2)^s d\mathbf{k}, \text{ where } \hat{v} := \mathcal{F}v \text{ for } v \in \mathcal{V}'. \quad (3.6)$$

The spatial Sobolev space $H^s(\mathbb{T}^2)$ can be alternatively identified using norm $\|\cdot\|_s$ and seminorm $|\cdot|_s$ (e.g., see Ruzhansky & Turunen (2009, Def. 3.2.2)) as follows for any $s \in \mathbb{R}$

$$\|f\|_s^2 := \sum_{\mathbf{k} \in \mathbb{Z}^2} (1 + |\mathbf{k}|^2)^s |\hat{f}(\mathbf{k})|^2, \text{ and } |f|_s^2 := \sum_{\mathbf{k} \in \mathbb{Z}_n^2 \setminus \{0\}} |\mathbf{k}|^{2s} |\hat{f}(\mathbf{k})|^2, \quad s \neq 0. \quad (3.7)$$

Moreover, we have the subsequent lemma in our specific case.

Theorem 3.2 (Functional norm “ \simeq ” negative norm). *If $f \in \mathcal{H} := L^2(\mathbb{T}^2)/\mathbb{R}$, $\|f\|_{\mathcal{H}'} = |f|_{-1}$.*

Spatially, we realize a regularized negative Sobolev norm by the Fast Fourier Transform (FFT):

$$\|f\|_{-1, \alpha, n}^2 := \sum_{\mathbf{k} \in \mathbb{Z}_n^2 \setminus \{0\}} (\alpha + |\mathbf{k}|^2)^{-1} |\hat{f}(\mathbf{k})|^2, \text{ where } \mathbb{Z}_n^2 := (\mathbb{Z}/n\mathbb{Z})^2, \text{ and } \alpha \geq 0. \quad (3.8)$$

With this norm at hand, (3.3) and the Bochner norm in (3.4) of the residual are realized to serve as the loss function in the fine-tuning

$$\eta_m(\mathbf{u}_{\mathcal{N}}, \partial_t \mathbf{u}_{\mathcal{N}}) := \|R(\mathbf{u}_{\mathcal{N}})(t_m, \cdot)\|_{-1, \alpha, n}. \quad (3.9)$$

Parseval identity. In the context of using an optimization algorithm to train an NN as a function approximator, it is known (e.g., Siegel et al. (2023)) that whether the integral in the loss function is accurately computed affects whether the NN can achieve the scientific computing level of accuracy. For example, on a uniform grid, the accuracy of the mesh-weighted spatial ℓ^2 -norm as the numerical quadrature is highly affected by the *local* smoothness of the function in consideration. Nevertheless, computing the integral in the frequency domain yields exponentially convergent approximations (Trefethen & Weideman, 2014, Theorem 3.1) thanks to the Parseval identity.

Why functional-type norm for the residual evaluation? We note that in “physics-informed” learning of operators, for example, in Li et al. (2024c), the PDE residual is evaluated using L^2 -norm as loss. In the meantime, *positive* Sobolev norm, which is local in terms of differential operators, is used in Li et al. (2022). To our best knowledge, the H^{-1} -functional norm has not been applied in either function learning or operator learning problems. In fact, the relation of the Gelfand triple is so simple and elegant, leading to Theorem 3.2, that the functional norm is nothing but an inverse frequency weight in the frequency domain that emphasizes the learning of *low frequency* information. This suits especially well for the learning problem of NSE. Quite contrary to the intuition of FNO variants having the frequency truncation that results error dominating in the high-frequency part, it has been discovered in Lippe et al. (2023) that the error of operator learning for NSE is still dominant in the lower end of the spectrum. This has been corroborated in our study as well, see Figure 8 and Figure 6. For a more detailed and mathematically enriched discussion on why functional norm is not widely popular in traditional numerical methods, we refer the reader to Appendix B.

3.3 NEW TRAINING-FINE-TUNING PARADIGM

With the trajectory-to-trajectory learner and the error estimators, we propose a new training-fine-tuning paradigm. Another important motivation of this new paradigm is that, experimentally, all FNO variants capture the statistical properties of the 2D turbulence (Benzi et al., 1987; Boffetta & Ecke, 2012) quickly in training. Even after a single epoch, the evaluations of ST-FNO converge to a tight neighborhood of the ground truth in terms of the frequency signature of the energy cascade of the flow (for details, please refer to Appendix D.3). Then, 95% of the FLOPs spent in training has marginal improvements, which motivates us to rethink a more efficient paradigm than end-to-end, in the meantime not needing to initiate the expensive training of another nonlinear corrector as PDE-refiner (Lippe et al., 2023) in the postprocessing phase.

The computation of extra fields. In evaluating the loss accurately and relying on this evaluation to apply the gradient method, another key barrier is that the extra field variables in evaluating the residual for the velocity (3.2), one needs to compute $\partial_t \mathbf{u}_{\mathcal{N}}$. Instead of a common approach of using NO to represent the extra field variables (Wen et al., 2022; Brandstetter et al., 2023), we opt

to apply a traditional numerical solver $G_\gamma(\cdot)$ with an extremely fine time step, e.g., $\mathcal{O}((\Delta t)^\gamma)$ for $\gamma \geq 2$, to compute these extra field variables (Line 8 of Algorithm 1), while preserving the computational graph for auto-differentiation. We note that, in NSE simulations using traditional numerical solvers, for efficiency, the magnitude of this fine time step is never realistic or attainable due to time marching. Given the training data with trajectories at $\{t_1, \dots, t_\ell\}$ aiming to predict the trajectories at $\{t_{\ell+1}, \dots, t_{\ell+n_t}\}$, the new paradigm to train and fine-tune ST-FNO is outlined in Algorithm 1.

Algorithm 1 The new parallel-in-time training-fine-tuning strategy in small data regime.

Input: ST-FNO $G_{\theta, \Theta} := \tilde{Q}_\theta \circ G_\Theta$; time stepping scheme $G_\alpha(\cdot)$; optimizer $\mathcal{D}(\theta, \nabla_\theta(\cdot))$; training dataset: solution trajectories at $[t_1, \dots, t_{\ell'}]$ as input and at $[t_{\ell'+1}, \dots, t_{\ell'+n_t}]$ as output.

- 1: Train the ST-FNO model until the energy signature matches the energy cascade.
- 2: **Freeze Θ in \mathcal{G}_Θ of ST-FNO $\mathcal{G}_{\theta, \Theta}$, keep \tilde{Q}_θ trainable.**
- 3: Cast all nn.Module involved and tensors to torch.float64 and torch.complex128 hereafter.

Input: Evaluation dataset: solution trajectories at $[t_1, \dots, t_\ell]$ as input, output time step n_t .

- 4: **for** $m = \ell, \dots, \ell + n_t - 1$ **do**
- 5: Extract the latent fields $\mathbf{v}_\mathcal{N}^{(m+1)}$ output of G_Θ at t_{m+1} and hold them fixed.
- 6: By construction of ST-FNO: $\mathbf{u}_\mathcal{N}^{(m+1)}(\theta) := \mathbf{u}_\mathcal{N}^{(\ell)} + \tilde{Q}_\theta(\mathbf{v}_\mathcal{N}^{(m+1)})$ for all m .
- 7: **for** $j = 1, \dots, \text{Iter}_{\max}$ **do**
- 8: March one step with $(\Delta t)^\gamma$ using G_γ and approximate the $\partial_t \mathbf{u}_\mathcal{N}$ as follows:
 $D_t \mathbf{u}_\mathcal{N}^{(m+1)}(\theta) := (\Delta t)^{-\gamma} (G_\gamma(\mathbf{u}_\mathcal{N}^{(m+1)}(\theta)) - \mathbf{u}_\mathcal{N}^{(m+1)}(\theta))$ for all m .
- 9: Compute $\eta^2 := \sum_m \eta_m^2(\mathbf{u}_\mathcal{N}^{(m+1)}(\theta), D_t \mathbf{u}_\mathcal{N}^{(m+1)}(\theta))$ for all evaluation time steps.
- 10: Apply the optimizer to update parameters in \tilde{Q} : $\theta \leftarrow \mathcal{D}(\theta, \nabla_\theta(\eta^2))$.
- 11: Forward pass only through \tilde{Q} to update $\mathbf{u}_\mathcal{N}^{(m+1)} \leftarrow \mathbf{u}_\mathcal{N}^{(\ell)} + \tilde{Q}_\theta(\mathbf{v}_\mathcal{N}^{(m+1)})$ for all m .

Output: A sequence of velocity profiles at corresponding time steps $\{\mathbf{u}_\mathcal{N}^{(m)}\}_{m=\ell+1}^{\ell+n_t}$.

Interpretations of the theoretical results. Theorem 3.1 establishes that the functional norm of the residual $R(\mathbf{u}_\mathcal{N})$, without accessing \mathbf{u} , is a good representation of the error $\mathbf{u} - \mathbf{u}_\mathcal{N}$ in Bochner norms. Estimate (3.4) indicates that reducing the *a posteriori* error estimator is a *necessary* condition for reducing the true error. While estimate (3.5) is more delicate in that it is only reliable when $\mathbf{u}_\mathcal{N}$ gets “close” to \mathbf{u} . Theorem 3.2 lays the foundation to accurately evaluate this functional norm via FFT. Nevertheless, (3.5) serves as a guideline to design the “refining” procedure (lines 11 in Algorithm 1), where the error estimator moves to refine the next time step once the error in the previous one becomes less than a given threshold.

4 NUMERICAL EXPERIMENTS

4.1 ILLUSTRATIVE EXAMPLE: TAYLOR-GREEN VORTEX

In this illustrative example, we examine the 2D Taylor-Green vortex model (Taylor & Green, 1937), whose analytical solution is *known* and frequently employed as a benchmark for evaluating traditional numerical schemes for NSE (Gottlieb & Orszag, 1977). We create a toy train dataset with 10 trajectories on a 256^2 grid with varying numbers of vortices per wavelength, and the test sample has an unseen number of vortices yet still can be resolved up to the Nyquist scale. For details please refer to Appendix C.2. The FNO3d used in this example is scaled down to 1 layer.

4.2 2D ISOTROPIC TURBULENCE

This meta-example contains a series of examples of isotropic turbulence (McWilliams, 1984), featured in various work such as Kochkov et al. (2021); Li et al. (2020). The energy and the enstrophy spectra satisfy the energy law of turbulence first proposed by A. N. Kolmogorov (Kolmogorov, 1941). The data are generated using a second-order time-stepping scheme that is proven in theory to preserve the inverse cascade of the energy/enstrophy spectra (Wang, 2012; Gottlieb et al., 2012). We consider two cases:

- (I) NSE benchmark in Li et al. (2020), $\nu = 10^{-3}$, $\omega_0 \sim \mathcal{N}(0, (-\Delta + \tau^2 I)^{-\alpha/2})$, the energy density in wavenumber $k := |\mathbf{k}|$ is $f(k) \sim (k^2 + \tau^2)^{-\alpha}$, no drag, fixed force with low wavenumber.
- (II) The famous decaying turbulence discovered by the McWilliams (McWilliams, 1984), the initial power spectrum $|\psi(k)|^2 \sim k^{-1}(\tau^2 + (k/k_0)^4)^{-1}$ is chosen that the decaying is slow and the enstrophy density evolves to the energy cascade of Kolmogorov flows featured in Kochkov et al. (2021); Lippe et al. (2023); Sun et al. (2023).

Tables 2 and 3 report the results for example (I) and example (II), respectively. Table 4 reports the computational efficiency comparison with roll-out and traditional solver.

5 CONCLUSION AND LIMITATIONS

We designed a new pipeline to train and fine-tune a new spatiotemporal modification of FNO to get close to the true solution (not the ground truth generated by the numerical solver) of NSE in the turbulent regime. The evaluation errors in benchmark problems are up to 10^5 times better than the non-fine-tuned FNO variants trained by a simple end-to-end pipeline. Due the exploitation of the intricate connections with traditional spectral methods, e.g., the optimally learned parameters correspond to a Fourier-Galerkin projection (Kovachki et al., 2021) using the Fourier basis, only FNO-based neural operators benefit from this advancement. How to generalize this new pipeline to other types of operator learners in Appendix B will be worthy of exploration.

Table 1: Results for Taylor-Green vortex $\varepsilon := \omega_{\text{true}} - \omega_{\mathcal{N}}$, the relative errors at the final time step.

	Evaluation after training		After fine-tuning	
	$\ \varepsilon\ _{L^2}$	$\ R\ _{-1,n}$	$\ \varepsilon\ _{L^2}$	$\ R\ _{-1,n}$
ST-FNO3d	1.94×10^{-1}	2.18×10^{-1}	1.24×10^{-7}	3.21×10^{-7}
PS+RK2 (GT)	5.91×10^{-6}	1.16×10^{-5}	N/A	N/A

Table 2: Results for forced turbulence, original example from Li et al. (2020). $\varepsilon := \omega_{\mathcal{S}} - \omega_{\mathcal{N}}$

	Evaluation after training		After fine-tuning	
	$\ \varepsilon\ _{L^2}$	$\ R\ _{-1,n}$	$\ \varepsilon\ _{L^2}$	$\ R\ _{-1,n}$
ST-FNO3d 10 ep + L^2 FT	2.08×10^{-2}	1.27×10^{-2}	2.82×10^{-4}	2.78×10^{-5}
ST-FNO3d 10 ep + H^{-1} FT	–	–	3.16×10^{-4}	4.59×10^{-7}

Table 3: Evaluation metrics of the McWilliams isotropic turbulence example. All models are trained using on 64×64 mesh and evaluated on 256×256 mesh. $\varepsilon := \mathbf{u}_{\mathcal{S}} - \mathbf{u}_{\mathcal{N}}$ or $\omega_{\mathcal{S}} - \omega_{\mathcal{N}}$. \mathcal{S} stands for the Fourier approximation space (3.1) on a 256×256 fine grid. $r := R(\mathbf{u}_{\mathcal{N}})$ or $R(\omega_{\mathcal{N}})$. $\mathcal{Y} := H^{-1}(\mathbb{T}^2)$. All error norms are evaluated at the final time step. H^{-1} appending model means the training uses the difference in the H^{-1} -norm as the loss function. Errors are measured for 32 trajectories in the test dataset. Fine-tuning uses 100 iterations of Adam optimizer in line 10.

	Evaluation after training		After fine-tuning	
	$\ \varepsilon\ _{L^2}$	$\ R\ _{-1,n}$	$\ \varepsilon\ _{L^2}$	$\ R\ _{-1,n}$
ST-FNO3d 10 epochs + H^{-1} train & H^{-1} FT	6.54×10^{-2}	6.19×10^{-2}	5.71×10^{-3}	2.24×10^{-5}
ST-FNO3d 10 epochs + L^2 train & FT	3.69×10^{-2}	2.35×10^{-2}	1.79×10^{-3}	9.55×10^{-5}
ST-FNO3d 10 epochs + L^2 train & H^{-1} FT	–	–	2.88×10^{-4}	4.02×10^{-6}

Table 4: The FLOPs/runtime comparison. FNO3d with a roll-out prediction (10 steps to predict next 10), ST-FNO3d (spatiotemporal prediction), and a traditional numerical solver’s time marching (IMEX 4-th order Runge-Kutta). Spatial grid size 256×256 , prediction time interval from $t = 4.5$ to $t = 5.5$ (stable energy cascade regime), numerical solver $\Delta t = 10^{-3}$, roll-out $\delta t = 10\Delta t$, ST-FNO3d $\delta t = 25\Delta t$. Note that ST-FNO3d gets all 40 steps in a single evaluation, the numerical solver has to march 1000 steps, and the roll-out prediction inference is supposed to march 10 times yet becomes de-correlated around $t \approx 4.8$ after the third step. **Runtimes are reported for a single instance averaging 20 runs with a 100w GPU power limiter on a single RTX 4500.**

	Architectures			GFLOPs			Runtime ($\times 10^{-2}$ s)		# params
	layers	channel	modes	Eval/Step	FT	All	Single eval	All	
FNO3d roll-out Example 2	4	10	(5, 32)	28.7	N/A	2.9M	4.1 ± 1.5	45.6 ± 6.7	16.38m
ST-FNO3d Example 2	$4 + \frac{1}{10}$	10	(5, 32)	1.06	2.51	337.6	22.7 ± 4.0	162 ± 4.1	16.42m
IMEX Runge-Kutta	N/A	N/A	N/A	3.98	N/A	4M	0.449 ± 0.3	450 ± 8.7	10

REFERENCES

- Mark Ainsworth and J Tinsley Oden. A unified approach to a posteriori error estimation using element residual methods. *Numerische Mathematik*, 65:23–50, 1993.
- Mark Ainsworth and J Tinsley Oden. A posteriori error estimation in finite element analysis. *Computer methods in applied mechanics and engineering*, 142(1-2):1–88, 1997.
- Kamyar Azizzadenesheli, Nikola Kovachki, Zongyi Li, Miguel Liu-Schiaffini, Jean Kossaifi, and Anima Anandkumar. Neural operators for accelerating scientific simulations and design. *Nature Reviews Physics*, pp. 1–9, 2024.
- Constantine A Balanis. *Advanced engineering electromagnetics*. John Wiley & Sons, 2012.
- Randolph E Bank and Bruno D Welfert. A posteriori error estimates for the Stokes problem. *SIAM Journal on Numerical Analysis*, 28(3):591–623, 1991.
- Francesca Bartolucci, Emmanuel de Bezenac, Bogdan Raonic, Roberto Molinaro, Siddhartha Mishra, and Rima Alaifari. Representation equivalent neural operators: a framework for alias-free operator learning. *Advances in Neural Information Processing Systems*, 36, 2024.
- R Benzi, S Patarnello, and P Santangelo. On the statistical properties of two-dimensional decaying turbulence. *Europhysics Letters*, 3(7):811, 1987.
- Christine Bernardi, Vivette Girault, and Yvon Maday. Mixed spectral element approximation of the Navier-Stokes equations in the stream-function and vorticity formulation. *IMA journal of numerical analysis*, 12(4):565–608, 1992.
- Guido Boffetta and Robert E Ecke. Two-dimensional turbulence. *Annual review of fluid mechanics*, 44:427–451, 2012.
- Daniele Boffi, Franco Brezzi, and Michel Fortin. *Mixed finite element methods and applications*, volume 44. Springer, 2013.
- Andrea Bonito, Claudio Canuto, Ricardo H Nochetto, and Andreas Veiser. Adaptive finite element methods. *arXiv preprint arXiv:2402.07273*, 2024.
- Nicolas Boullé and Alex Townsend. A mathematical guide to operator learning. *arXiv preprint arXiv:2312.14688*, 2023.
- Johannes Brandstetter, Rianne van den Berg, Max Welling, and Jayesh K Gupta. Clifford neural layers for PDE modeling. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=okwxL_c4x84.
- Steven L Brunton, Joshua L Proctor, and J Nathan Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the national academy of sciences*, 113(15):3932–3937, 2016.
- John C Butcher. On the attainable order of runge-kutta methods. *Mathematics of computation*, 19(91):408–417, 1965.
- Claudio Canuto, M Yousuff Hussaini, Alfio Quarteroni, and Thomas A Zang. *Spectral methods: evolution to complex geometries and applications to fluid dynamics*. Springer Science & Business Media, 2007.
- Lianghao Cao, Thomas O’Leary-Roseberry, Prashant K Jha, J Tinsley Oden, and Omar Ghattas. Residual-based error correction for neural operator accelerated infinite-dimensional bayesian inverse problems. *Journal of Computational Physics*, 486:112104, 2023.
- Shuhao Cao. Choose a Transformer: Fourier or Galerkin. *Advances in neural information processing systems*, 34:24924–24940, 2021.
- Carsten Carstensen and Stefan A Funken. Fully reliable localized error control in the fem. *SIAM Journal on Scientific Computing*, 21(4):1465–1484, 1999.

- Kathleen Champion, Bethany Lusch, J Nathan Kutz, and Steven L Brunton. Data-driven discovery of coordinates and governing equations. *Proceedings of the National Academy of Sciences*, 116(45): 22445–22451, 2019.
- Yanlai Chen and Shawn Koohy. GPT-PINN: Generative pre-trained physics-informed neural networks toward non-intrusive meta-learning of parametric pdes. *Finite Elements in Analysis and Design*, 228:104047, 2024.
- Alexandre Joel Chorin. Numerical solution of the Navier-Stokes equations. *Mathematics of computation*, 22(104):745–762, 1968.
- Maarten V de Hoop, Nikola B Kovachki, Nicholas H Nelsen, and Andrew M Stuart. Convergence rates for learning linear operators from noisy data. *SIAM/ASA Journal on Uncertainty Quantification*, 11(2):480–513, 2023.
- Gideon Dresdner, Dmitrii Kochkov, Peter Christian Norgaard, Leonardo Zepeda-Nunez, Jamie Smith, Michael Brenner, and Stephan Hoyer. Learning to correct spectral methods for simulating turbulent flows. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=wNBARGxoJn>.
- Yiheng Du, Nithin Chalapathi, and Aditi S. Krishnapriyan. Neural spectral methods: Self-supervised learning in the spectral domain. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=2DbVeuoa6a>.
- Alexandre Ern and Martin Vohralík. A posteriori error estimation based on potential and flux reconstruction for the heat equation. *SIAM Journal on Numerical Analysis*, 48(1):198–223, 2010.
- Lawrence C Evans. *Partial differential equations*, volume 19. American Mathematical Society, 2022.
- Vladimir Fanaskov, Alexander Rudikov, and Ivan Oseledets. Neural functional a posteriori error estimates. *arXiv preprint arXiv:2402.05585*, 2024.
- Julian Fischer. A posteriori modeling error estimates for the assumption of perfect incompressibility in the Navier–Stokes equation. *SIAM Journal on Numerical Analysis*, 53(5):2178–2205, 2015.
- Antonio Henrique De Oliveira Fonseca, Emanuele Zappala, Josue Ortega Caro, and David Van Dijk. Continuous spatiotemporal Transformer. In *International Conference on Machine Learning*, pp. 7343–7365. PMLR, 2023.
- I.M. Gel’fand and G.E. Shilov. *Generalized Functions, Volume 2: Spaces of Fundamental and Generalized Functions*. AMS Chelsea Publishing. American Mathematical Society, 2016. ISBN 9781470426590. URL <https://books.google.com/books?id=HPDdCwAAQBAJ>.
- Andrew Gillette, Brendan Keith, and Socratis Petrides. Learning robust marking policies for adaptive mesh refinement. *SIAM Journal on Scientific Computing*, 46(1):A264–A289, 2024.
- Vivette Girault and Pierre-Arnaud Raviart. *Finite element methods for Navier-Stokes equations: theory and algorithms*, volume 5. Springer Science & Business Media, 2012.
- Jonathan Goodman, Thomas Hou, and Eitan Tadmor. On the stability of the unsmoothed fourier method for hyperbolic equations. *Numerische Mathematik*, 67(1):93–129, 1994.
- Somdatta Goswami, Aniruddha Bora, Yue Yu, and George Em Karniadakis. Physics-informed deep neural operator networks. In *Machine Learning in Modeling and Simulation: Methods and Applications*, pp. 219–254. Springer, 2023.
- David Gottlieb and Steven A Orszag. *Numerical analysis of spectral methods: theory and applications*. SIAM, 1977.
- Sigal Gottlieb, Florentina Tone, Cheng Wang, Xiaoming Wang, and Djoko Wirosoetisno. Long time stability of a classical efficient scheme for two-dimensional Navier–Stokes equations. *SIAM Journal on Numerical Analysis*, 50(1):126–150, 2012.

- Daniel Greenfeld, Meirav Galun, Ronen Basri, Irad Yavneh, and Ron Kimmel. Learning to optimize multigrid pde solvers. In *International Conference on Machine Learning*, pp. 2415–2423. PMLR, 2019.
- Albert Gu, Tri Dao, Stefano Ermon, Atri Rudra, and Christopher Ré. Hippo: Recurrent memory with optimal polynomial projections. *Advances in neural information processing systems*, 33: 1474–1487, 2020.
- Gaurav Gupta, Xiongye Xiao, and Paul Bogdan. Multiwavelet-based operator learning for differential equations. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=LZDiWaC9CGL>.
- Jayesh K Gupta and Johannes Brandstetter. Towards multi-spatiotemporal-scale generalized PDE modeling. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=dPSTDbGtBY>.
- Jiequn Han, Arnulf Jentzen, and Weinan E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.
- Zhongkai Hao, Zhengyi Wang, Hang Su, Chengyang Ying, Yinpeng Dong, Songming Liu, Ze Cheng, Jian Song, and Jun Zhu. Gnot: A general neural operator Transformer for operator learning. In *International Conference on Machine Learning*, pp. 12556–12569. PMLR, 2023.
- Juncai He, Xinliang Liu, and Jinchao Xu. MgNO: Efficient parameterization of linear operators via multigrid. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=80xL034uEr>.
- Yinnian He and Weiwei Sun. Stability and convergence of the Crank–Nicolson/Adams–Bashforth scheme for the time-dependent Navier–Stokes equations. *SIAM Journal on Numerical Analysis*, 45(2):837–869, 2007.
- Jan S Hesthaven, Gianluigi Rozza, Benjamin Stamm, et al. *Certified reduced basis methods for parametrized partial differential equations*, volume 590. Springer, 2016.
- John G Heywood and Rolf Rannacher. Finite element approximation of the nonstationary Navier–Stokes problem, part ii: stability of solutions and error estimates uniform in time. *SIAM journal on numerical analysis*, 23(4):750–777, 1986.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf.
- Maarten V. de Hoop, Daniel Zhengyu Huang, Elizabeth Qian null, and Andrew M. Stuart. The cost-accuracy trade-off in operator learning with neural networks. *Journal of Machine Learning*, 1(3):299–341, June 2022. ISSN 2790-2048. doi: 10.4208/jml.220509. URL <http://dx.doi.org/10.4208/jml.220509>.
- Thomas Y Hou. Blow-up or no blow-up? a unified computational and analytic approach to 3d incompressible euler and Navier–Stokes equations. *Acta Numerica*, 18:277–346, 2009.
- Jun Hu and Pengzhan Jin. A hybrid iterative method based on mionet for pdes: Theory and numerical examples. *arXiv preprint arXiv:2402.07156*, 2024.
- Ru Huang, Ruipeng Li, and Yuanzhe Xi. Learning optimal multigrid smoothers via neural networks. *SIAM Journal on Scientific Computing*, 45(3):S199–S225, 2022.
- Ru Huang, Kai Chang, Huan He, Ruipeng Li, and Yuanzhe Xi. Reducing operator complexity in algebraic multigrid with machine learning approaches. *arXiv preprint arXiv:2307.07695*, 2023.
- Haris Iqbal. Harisiqbal88/plotneuralnet v1.0.0, 2018. URL <https://zenodo.org/record/2526396>.

- Hans Johnston and Jian-Guo Liu. Accurate, stable and efficient Navier–Stokes solvers based on explicit treatment of the pressure term. *Journal of Computational Physics*, 199(1):221–259, 2004.
- George Em Karniadakis, Moshe Israeli, and Steven A Orszag. High-order splitting methods for the incompressible Navier-Stokes equations. *Journal of computational physics*, 97(2):414–443, 1991.
- Aly-Khan Kassam and Lloyd N Trefethen. Fourth-order time-stepping for stiff pdes. *SIAM Journal on Scientific Computing*, 26(4):1214–1233, 2005.
- Ryan Keisler. Forecasting global weather with graph neural networks. *arXiv preprint arXiv:2202.07575*, 2022.
- Georgios Kissas, Jacob H Seidman, Leonardo Ferreira Guilhoto, Victor M Preciado, George J Pappas, and Paris Perdikaris. Learning operators with coupled attention. *Journal of Machine Learning Research*, 23(215):1–63, 2022.
- Dmitrii Kochkov, Jamie A Smith, Ayya Alieva, Qing Wang, Michael P Brenner, and Stephan Hoyer. Machine learning–accelerated computational fluid dynamics. *Proceedings of the National Academy of Sciences*, 118(21):e2101784118, 2021.
- Andrey Nikolaevich Kolmogorov. The local structure of turbulence in incompressible viscous fluid for very large reynolds. *Numbers. In Dokl. Akad. Nauk SSSR*, 30:301, 1941.
- Nikola Kovachki, Samuel Lanthaler, and Siddhartha Mishra. On universal approximation and error bounds for Fourier neural operators. *Journal of Machine Learning Research*, 22(290):1–76, 2021.
- Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with applications to pdes. *Journal of Machine Learning Research*, 24(89):1–97, 2023.
- Heinz-Otto Kreiss and Joseph Oliger. Stability of the fourier method. *SIAM Journal on Numerical Analysis*, 16(3):421–433, 1979.
- Hwar C Ku, Thomas D Taylor, and Richard S Hirsh. Pseudospectral methods for solution of the incompressible Navier-Stokes equations. *Computers & fluids*, 15(2):195–214, 1987.
- Samuel Lanthaler, Roberto Molinaro, Patrik Hadorn, and Siddhartha Mishra. Nonlinear reconstruction for operator learning of PDEs with discontinuities. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=CrfhZAsJDsZ>.
- Peter D Lax. *Functional analysis*, volume 55. John Wiley & Sons, 2002.
- Sanjiva K Lele. Compact finite difference schemes with spectral-like resolution. *Journal of computational physics*, 103(1):16–42, 1992.
- Zhijie Li, Wenhui Peng, Zelong Yuan, and Jianchun Wang. Long-term predictions of turbulence by implicit u-net enhanced fourier neural operator. *Physics of Fluids*, 35(7), 2023a.
- Zijie Li, Kazem Meidani, and Amir Barati Farimani. Transformer for partial differential equations’ operator learning. *Transactions on Machine Learning Research*, 2023b. ISSN 2835-8856. URL <https://openreview.net/forum?id=EPPqt3uERT>.
- Zijie Li, Dule Shu, and Amir Barati Farimani. Scalable transformer for pde surrogate modeling. *Advances in Neural Information Processing Systems*, 36, 2024a.
- 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.
- Zongyi Li, Miguel Liu-Schiaffini, Nikola Borislavov Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Learning chaotic dynamics in dissipative systems. In *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=1C36tFZn7sR>.

- Zongyi Li, Nikola Kovachki, Chris Choy, Boyi Li, Jean Kossaifi, Shourya Ota, Mohammad Amin Nabian, Maximilian Stadler, Christian Hundt, Kamyar Azizzadenesheli, et al. Geometry-informed neural operator for large-scale 3d pdes. *Advances in Neural Information Processing Systems*, 36, 2024b.
- 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*, 1(3):1–27, 2024c.
- Marten Lienen, Jan Hansen-Palmus, David Lüdke, and Stephan Günnemann. Generative diffusion for 3d turbulent flows. *arXiv preprint arXiv:2306.01776*, 2023.
- Jacques Louis Lions and Enrico Magenes. *Non-homogeneous boundary value problems and applications: Vol. 1*, volume 181. Springer Science & Business Media, 2012.
- Phillip Lippe, Bastiaan S. Veeling, Paris Perdikaris, Richard E Turner, and Johannes Brandstetter. PDE-refiner: Achieving accurate long rollouts with neural PDE solvers. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=Qv64681lWS>.
- Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via deepnet based on the universal approximation theorem of operators. *Nature machine intelligence*, 3(3):218–229, 2021.
- Martine Marion and R Temam. Nonlinear Galerkin methods: the finite elements case. *Numerische Mathematik*, 57:205–226, 1990.
- Nick McGreivy and Ammar Hakim. Invariant preservation in machine learned pde solvers via error correction. *arXiv preprint arXiv:2303.16110*, 2023.
- James C McWilliams. The emergence of isolated coherent vortices in turbulent flow. *Journal of Fluid Mechanics*, 146:21–43, 1984.
- J Tinsley Oden, Weihai Wu, and Mark Ainsworth. An a posteriori error estimate for finite element approximations of the Navier-Stokes equations. *Computer Methods in Applied Mechanics and Engineering*, 111(1-2):185–202, 1994.
- Steven A. Orszag. Numerical simulation of incompressible flows within simple boundaries. i. Galerkin (spectral) representations. *Studies in Applied Mathematics*, 50(4):293–327, 1971a. ISSN 1467-9590. doi: 10.1002/sapm1971504293. URL <http://dx.doi.org/10.1002/sapm1971504293>.
- Steven A. Orszag. On the elimination of aliasing in finite-difference schemes by filtering high-wavenumber components. *Journal of the Atmospheric Sciences*, 28(6):1074–1074, 1971b. ISSN 1520-0469. doi: 10.1175/1520-0469(1971)028<1074:oteoi>2.0.co;2. URL [http://dx.doi.org/10.1175/1520-0469\(1971\)028<1074:OTE0AI>2.0.CO;2](http://dx.doi.org/10.1175/1520-0469(1971)028<1074:OTE0AI>2.0.CO;2).
- Steven A Orszag. Comparison of pseudospectral and spectral approximation. *Studies in Applied Mathematics*, 51(3):253–259, 1972.
- Anthony T Patera, Gianluigi Rozza, et al. Reduced basis approximation and a posteriori error estimation for parametrized partial differential equations, 2007.
- GS Patterson and Steven A Orszag. Spectral calculations of isotropic turbulence: Efficient removal of aliasing interactions. *Physics of Fluids*, 14(11):2538, 1971.
- Jaak Peetre. On spaces of Triebel–Lizorkin type. *Arkiv för Matematik*, 13(1):123–130, 1975.
- Kai Qi and Jian Sun. Gabor-filtered fourier neural operator for solving partial differential equations. *Computers & Fluids*, 274:106239, 2024.
- Md Ashiqur Rahman, Zachary E Ross, and Kamyar Azizzadenesheli. U-NO: U-shaped neural operators. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=j3oQF9coJd>.

- Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019.
- Rolf Rannacher. *Finite element methods for the incompressible Navier-Stokes equations*. Springer, 2000.
- Bogdan Raonic, Roberto Molinaro, Tim De Ryck, Tobias Rohner, Francesca Bartolucci, Rima Alaifari, Siddhartha Mishra, and Emmanuel de Bézenac. Convolutional neural operators for robust and accurate learning of pdes. *Advances in Neural Information Processing Systems*, 36, 2024.
- Sergey Repin. *A Posteriori Estimates for Partial Differential Equations*. Walter de Gruyter, September 2008. ISBN 9783110191530. URL <http://dx.doi.org/10.1515/9783110203042>.
- Sergey I Repin. A posteriori error estimation for nonlinear variational problems by duality theory. *Journal of Mathematical Sciences*, 99:927–935, 2000.
- Michael Ruzhansky and Ville Turunen. *Pseudo-differential operators and symmetries: background analysis and advanced topics*, volume 2. Springer Science & Business Media, 2009.
- Jie Shen. Efficient spectral-galerkin method i. direct solvers of second-and fourth-order equations using legendre polynomials. *SIAM Journal on Scientific Computing*, 15(6):1489–1505, 1994.
- Dule Shu, Zijie Li, and Amir Barati Farimani. A physics-informed diffusion model for high-fidelity flow field reconstruction. *Journal of Computational Physics*, 478:111972, 2023.
- Jonathan W. Siegel, Qingguo Hong, Xianlin Jin, Wenrui Hao, and Jinchao Xu. Greedy training algorithms for neural networks and applications to pdes. *Journal of Computational Physics*, 484:112084, 2023. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2023.112084>. URL <https://www.sciencedirect.com/science/article/pii/S0021999123001791>.
- Leslie N Smith and Nicholay Topin. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial intelligence and machine learning for multi-domain operations applications*, volume 11006, pp. 369–386. SPIE, 2019.
- Zhiqing Sun, Yiming Yang, and Shinjae Yoo. A neural pde solver with temporal stencil modeling. In *International Conference on Machine Learning*, pp. 33135–33155. PMLR, 2023.
- Eitan Tadmor. Stability analysis of finite difference, pseudospectral and fourier–galerkin approximations for time-dependent problems. *SIAM review*, 29(4):525–555, 1987.
- Ali Taghibakhshi, Scott MacLachlan, Luke Olson, and Matthew West. Optimization-based algebraic multigrid coarsening using reinforcement learning. *Advances in neural information processing systems*, 34:12129–12140, 2021.
- Ali Taghibakhshi, Nicolas Nytko, Tareq Uz Zaman, Scott MacLachlan, Luke Olson, and Matthew West. Mg-gnn: Multigrid graph neural networks for learning multilevel domain decomposition methods. In *International Conference on Machine Learning*, pp. 33381–33395. PMLR, 2023.
- Geoffrey Ingram Taylor and Albert Edward Green. Mechanism of the production of small eddies from large ones. *Proceedings of the Royal Society of London. Series A-Mathematical and Physical Sciences*, 158(895):499–521, 1937.
- Roger Temam. *Navier–Stokes equations and nonlinear functional analysis*. SIAM, 1995.
- Alasdair Tran, Alexander Mathews, Lexing Xie, and Cheng Soon Ong. Factorized fourier neural operators. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=tmIiMPL4IPa>.
- Lloyd N Trefethen and JAC Weideman. The exponentially convergent trapezoidal rule. *SIAM review*, 56(3):385–458, 2014.
- Kiwon Um, Robert Brand, Yun Raymond Fei, Philipp Holl, and Nils Thuerey. Solver-in-the-loop: Learning from differentiable physics to interact with iterative pde-solvers. *Advances in Neural Information Processing Systems*, 33:6111–6122, 2020.

- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Andreas Veiser and Rüdiger Verfürth. Explicit upper bounds for dual norms of residuals. *SIAM journal on numerical analysis*, 47(3):2387–2405, 2009.
- Andreas Veiser and Rüdiger Verfürth. Poincaré constants for finite element stars. *IMA Journal of Numerical Analysis*, 32(1):30–47, 2012.
- Rüdiger Verfürth. A posteriori error estimators for the Stokes equations. *Numerische Mathematik*, 55(3):309–325, 1989.
- Rüdiger Verfürth. A posteriori error estimates for nonlinear problems. finite element discretizations of elliptic equations. *Mathematics of Computation*, 62(206):445–475, 1994.
- Rüdiger Verfürth. A posteriori error estimates for finite element discretizations of the heat equation. *Calcolo*, 40(3):195–212, 2003.
- Rüdiger Verfürth. *A posteriori error estimation techniques for finite element methods*. OUP Oxford, 2013.
- Sifan Wang, Hanwen Wang, and Paris Perdikaris. Learning the solution operator of parametric partial differential equations with physics-informed deeponets. *Science advances*, 7(40):eabi8605, 2021.
- Xiaoming Wang. An efficient second order in time scheme for approximating long time statistical properties of the two dimensional Navier–Stokes equations. *Numerische Mathematik*, 121(4):753–779, 2012.
- E Weinan and Jian-Guo Liu. Projection method i: convergence and numerical boundary layers. *SIAM journal on numerical analysis*, pp. 1017–1057, 1995.
- Gege Wen, Zongyi Li, Kamyar Azizzadenesheli, Anima Anandkumar, and Sally M Benson. U-fno—an enhanced fourier neural operator-based deep-learning model for multiphase flow. *Advances in Water Resources*, 163:104180, 2022.
- Liang Yan and Tao Zhou. An adaptive surrogate modeling based on deep neural networks for large-scale bayesian inverse problems. *Communications in Computational Physics*, 28(5):2180–2205, June 2020. doi: 10.4208/cicp.oa-2020-0186.
- Jiachen Yang, Tarik Dzanic, Brenden Petersen, Jun Kudo, Ketan Mittal, Vladimir Tomov, Jean-Sylvain Camier, Tuo Zhao, Hongyuan Zha, Tzanio Kolev, et al. Reinforcement learning for adaptive mesh refinement. In *International Conference on Artificial Intelligence and Statistics*, pp. 5997–6014. PMLR, 2023.
- Hongkai Zheng, Weili Nie, Arash Vahdat, Kamyar Azizzadenesheli, and Anima Anandkumar. Fast sampling of diffusion models via operator learning. In *International Conference on Machine Learning*, pp. 42390–42402. PMLR, 2023.

A NOTATIONS

Table 5: Notations used in an approximate chronological order and their meaning in this work.

Notation	Meaning
Ω	$\Omega = (0, 2\pi)^2$ or $(0, 1)^2$ the modeling domain in \mathbb{R}^2
\mathcal{V}, \mathcal{H}	Hilbert spaces defined on the domain Ω above, $f \in \mathcal{H} : \Omega \rightarrow \mathbb{R}$
\mathcal{X}, \mathcal{Y}	product spaces $\prod_{j \in \Lambda} \mathcal{H}$
$\ $	concatenation, for $f_j \in \mathcal{H}$, $\ _{j \in \Lambda} f_j \in \prod_{j \in \Lambda} \mathcal{H}$
ν	viscosity, the inverse of the Reynolds number, strength of diffusion
\mathbb{T}^2	$\mathbb{T}^2 := S^1 \times S^1$ where S^1 is homeomorphic to $[0, 1)$ with the point 1 being 0.
$H^1(\mathbb{T}^2)$	$\{u \in H^1(\Omega) : u \text{ satisfies the periodical boundary condition}\}$
$\mathbf{H}^1(\mathbb{T}^2)$	$\mathbf{H}^1(\mathbb{T}^2) := H^1(\mathbb{T}^2) \times H^1(\mathbb{T}^2)$
$L^p(\mathcal{T}; \mathcal{V})$	the Bochner spaces containing $\{u : \mathcal{T} \rightarrow \mathcal{V} \mid \int_{\mathcal{T}} \ u(t, \cdot)\ _{\mathcal{V}}^p dt < +\infty\}$, $p \in [1, \infty)$
$L^\infty(\mathcal{T}; \mathcal{V})$	Bochner space containing $\{u : \mathcal{T} \rightarrow \mathcal{H} \mid \text{ess sup}_{t \in \mathcal{T}} \ u(t, \cdot)\ _{\mathcal{V}} < +\infty\}$
\mathbf{u}	the true weak solution to the NSE $\mathbf{u}(t, \cdot) \in \mathcal{V}$ for any t
$\mathbf{u}_S^{(l)}$	the ground truth data generated at t_l by the numerical solver in $S \subset \mathcal{V}$
$\mathbf{u}_{\mathcal{N}}^{(m)}$	neural operator evaluations at t_m that can be embedded in $S \subset \mathcal{V}$
(u, v) or (\mathbf{u}, \mathbf{v})	the L^2 -inner product on \mathcal{H} , $(u, v) := \int_{\Omega} uv \, dx$
$\ u\ _s$	the H^s -norm of u , computed by (3.7), $\ u\ := \ u\ _0$ falls back to L^2 -norm
$ u _s$	the H^s -seminorm of u , computed by (3.7), is a norm on $H^s(\mathbb{T}^2)/\mathbb{R}$
\lesssim	$a \lesssim b$ means that $\exists c$ independent of asymptotics if any such that $a \leq cb$
\simeq	$a \simeq b \Leftrightarrow a \lesssim b$ and $b \lesssim a$
\mathcal{V}'	dual space of \mathcal{V} , contains all continuous functionals f such that $ f(v) \lesssim \ v\ _{\mathcal{V}}$
$\langle f, v \rangle$	the pairing between $v \in \mathcal{V}$ and $f \in \mathcal{V}'$, can be identified as (f, v) if $f \in \mathcal{H}$
$\mathcal{V} \hookrightarrow \mathcal{H}$	\mathcal{V} is continuously embedded in \mathcal{H} such that $\ u\ _{\mathcal{H}} \lesssim \ u\ _{\mathcal{V}}$
$\mathcal{V} \Subset \mathcal{H} \Subset \mathcal{V}'$	compact embeddings by Poincaré inequality, $\mathcal{H} = L^2(\mathbb{T}^2)$, $\mathcal{V} = H^1(\mathbb{T}^2)$
$\mathbf{A} : \mathbf{B}$	$\mathbf{A} : \mathbf{B} = \sum_{1 \leq i, j \leq 2} A_{ij} B_{ij}$ for $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{2 \times 2}$
$\mathbf{a} \otimes \mathbf{b}$	$\mathbf{a}\mathbf{b}^\top$ if both are viewed as column vectors

B DETAILED LITERATURE REVIEW AND MOTIVATIONS

Interplay of deep learning and PDEs. PDE solvers are function learners to represent PDE solutions using neural networks Han et al. (2018); Raissi et al. (2019); Chen & Koohy (2024). PDE discovery encompasses all the techniques dedicated to identifying and optimizing PDE coefficients from data Brunton et al. (2016); Champion et al. (2019). Recently, hybrid solver approach has been explored in Greenfeld et al. (2019); Kochkov et al. (2021); Huang et al. (2023); Taghibakhshi et al. (2023, 2021); Hu & Jin (2024); Huang et al. (2022). Reinforcement learning has been applied in the field of mesh generation to build more efficient solving pipelines Yang et al. (2023); Gillette et al. (2024). For neural operators, remarkable outcomes are achieved in diverse applications, such as weather forecasting Keisler (2022) and turbulent fluids simulations Shu et al. (2023); Li et al. (2023a); Lienen et al. (2023), the methods based on neural operators have significantly influenced the progress of the interplay between PDE and deep learning. This success was a natural outcome of several improvements brought to the field, for example, fast solution evaluations, a feature very appealing in many engineering applications Zheng et al. (2023); and the ability to provide mesh-free and resolution-independent solvers in cases of irregular domains Hao et al. (2023); Li et al. (2024b).

Neural operators. Looking at neural operator architectures, we can identify two different approaches. The first one creates “basis” (or frames) through nonlinear approximators and aggregates them linearly. Developed architectures that belong to this category include Deep Operator Network (DeepONet) Lu et al. (2021); Goswami et al. (2023); Wang et al. (2021), wherein aggregation occurs via linear combination. Similarly, Fourier Neural Operator (FNO) Li et al. (2020), and some of its variants Gupta et al. (2021); Rahman et al. (2023); Tran et al. (2023); Li et al. (2024b); Wen et al. (2022), aggregate different latent representations through convolution with a learnable kernel in the frequency domain. Additional examples with a linear aggregation with a U-Net meta-architecture

can be found in Raonic et al. (2024); He et al. (2024). Bartolucci et al. (2024) further explores the error analyses for the linear aggregations through the lens of the frame(let) theory. The second approach to designing neural operator architectures aims to obtain the “basis” through linear projections of the latent representations. In this scenario, the aggregation is nonlinear, for example, using a signal-dependent kernel integral. In this category, the most notable example is the scaled dot-product attention operator in Transformer Vaswani et al. (2017), which builds an instance-dependent kernel. In the context of operator learning applications, Transformer-based neural operators have been studied in Kissas et al. (2022); Cao (2021); Li et al. (2023b); Hao et al. (2023); Li et al. (2024a); Fonseca et al. (2023). It is also shown in Lanthaler et al. (2023) that nonlinear aggregations outperform its linear counterparts in learning solutions with less regularity.

Numerical methods for NSE. The NSE falls into the category of a stiff PDE (system) $\partial_t u = Lu + N(u) + f$, where f is the external forcing, L and $N(\cdot)$ are a linear and a nonlinear operator, respectively. Trefethen noted back in Kassam & Trefethen (2005) on the difficulty to design a time-stepping scheme as $N(\cdot)$ and L have to be treated differently. There are a long history of numerical methods for NSE we draw inspiration from. Petrov-Galerkin methods have been developed for NSE in Boffi et al. (2013); Girault & Raviart (2012). Nonlinear Galerkin method Marion & Temam (1990) inspires us to prove Theorem E.9. Chorin (1968) uses a clever trick to impose the divergence free condition without constructing a divergence-conforming finite element subspace. Shen (1994) designed various Galerkin methods in the space of orthogonal polynomials. Pioneered by Chorin, Shen, and E, projection methods are among the most popular schemes to solve NSE (see e.g., Weinan & Liu (1995) for a summary), also serves as an inspiration to add the Helmholtz layer. Bernardi et al. (1992) proposed a mixed discretization for the vorticity-streamfunction formulation.

Consistency-stability trade-offs. In view of the Lax equivalence principle (“consistency” + “stability” \implies “convergence” (Lax, 2002, Theorem 8)), the improvement of the stability of the method has a trade-off with a method’s consistency at the cost of the approximation capacity. Numerical methods for NSE is the epitome for such trade-off. For example, high-order explicit time stepping schemes offer better local truncation error estimates near boundary layers of the flow Lele (1992), yet the lack of stability is more severe and needs higher-order temporal smoothing. On the other hand, implicit schemes can be unconditionally stable for stiff or highly transient NSE with relatively large time steps $\mathcal{O}(1)$. The stability in implicit schemes becomes much less stringent on the time step, as no CFL condition is required. However, the solution at the next time step requires solving a linear or nonlinear system. Thus, computationally implicit schemes are usually an order of magnitude more expensive compared to explicit schemes.

De-aliasing filter sacrifices consistency for stability. One famous example of the consistency-stability trade-off is the 3/2-rule (also known as 2/3-dealiasing dealiasing filter) for the nonlinear convective term Orszag (1971b); Patterson & Orszag (1971); Hou (2009); Gottlieb & Orszag (1977) for pseudo-spectral methods Orszag (1971a; 1972). The highest 1/3 modes, the inclusion of which contributes to better approximation capacity, are filtered out to ensure long-term stability. Compromises such as the CFL condition and the 3/2-rule *must* be made for traditional numerical schemes to be stable, keeping the balance between stability and accuracy. These constraints apply to traditional numerical methods because any solver has to march a consecutive multitude of time steps, which makes the error propagation operator’s norm a product of many. Numerical results suggest that for pseudo-spectral spatial discretization with no higher-order Fourier smoothing temporally, the dealiasing filter is indispensable Tadmor (1987), as the time marching may experience numerical instability Kreiss & Oliger (1979); Goodman et al. (1994) without it. This is due to the nonlinear interaction in the convective term, which is caused by the amplification of high-frequency “aliasing” errors when the underlying solution lacks sufficient smoothness. In this study, the hybrid approach we adopted combines the strengths of NOs and traditional numerical solvers. There is no time marching consecutively for a multitude of time steps, which renders the method free of the stability constraints such as the CFL condition and the 3/2-rule.

Why and why not functional-type a posteriori error estimation? We consider the following enlightening Pythagorean-type identity for the true solution u satisfying $A(u) = f$

$$\|u - u_S\|_a^2 + \|u_S - u_N\|_a^2 = \|u - u_N\|_a^2 \quad \text{since } (u - u_S) \perp_a (u_S - u_N) \in \mathcal{V}. \quad (\text{B.1})$$

For the sake of presentation, we make some handy assumptions to get this identity: (1) $A : \mathcal{V} \rightarrow \mathcal{V}'$ denotes a linear operator on the Gelfand triple; (2) u_S is obtained through Galerkin methods [Evans \(2022\)](#), which is a projection onto a finite-dimensional approximation subspace of $\mathcal{S} \subset \mathcal{V}$; and (3) an evaluation of u_N can be continuously embedded into \mathcal{S} . Through the Riesz representation theorem, u_S is solved through $(u - u_S, v)_a := \langle A(u) - A(u_S), v \rangle = 0$, for any $v \in \mathcal{S}$. This bilinear form $(\cdot, \cdot)_a$ induces a (semi)norm $\|\cdot\|_a$ inheriting the topology from \mathcal{V} . Given this orthogonality, minimizing the difference between u_N and u_S becomes fruitless if $\|u - u_S\|_a$ is relatively big in the first place, and unnecessary computational resources may have been spent to get closer to u_S . Rather, [\(B.1\)](#) indicates that it is more efficient if one can design a method to reduce the error of $\|u - u_N\|_a$ directly, while circumventing the fact that u is not accessible.

Speaking of the *a posteriori* error estimation in traditional PDE discretization, part of the goal is to help the adaptive mesh refinement to get a better local basis. In this regard, the global error functional in negative Sobolev spaces must be approximated using localized L^2 residuals to indicate where the mesh needs to be refined. There are various compromises for this H^{-1} -to- L^2 representation to happen that renders the estimate inaccurate, such as discrete Poincaré constant [Veese & Verfürth \(2012\)](#) or inverse inequalities [Carstensen & Funken \(1999\)](#); [Veese & Verfürth \(2009\)](#), see also [\(Verfürth, 2013, §1.6.2\)](#). Functional-type *a posteriori* error estimation [Repin \(2008\)](#) consider the error as a functional, which is equivalent to error to bilinear form-associated norm as follows

$$R(u_S) \in \mathcal{V}', \text{ and } \langle R(u_S), v \rangle = (u - u_S, v)_a,$$

where the weak form of the PDE is

$$(u, v)_a = f(v) \quad \forall v \in \mathcal{V} \text{ and } (u_S, v)_a = f(v) \quad \forall v \in \mathcal{S} \subset \mathcal{V}.$$

The common approach is using the help from extra “flux” or “stress” dual variables ([\(Repin, 2008, §6.4\)](#), see also [Repin \(2000\)](#)) for how to get an accurate representation of the error functional in L^2 . For example, for the Stokes flow (steady-state viscous fluid, letting $\partial_t \mathbf{u} = 0$ and no nonlinear convection in [\(2.2\)](#), $\nu \sim \mathcal{O}(1)$), [\(Repin, 2008, §6.2\)](#) estimates error of an $\mathbf{H}(\text{div})$ - L^2 mixed discretization as follows

$$\nu \|\nabla(\mathbf{u} - \mathbf{u}_S)\| \leq \|\boldsymbol{\sigma} + q\mathbf{I} - \nu \nabla \mathbf{u}_S\| + C_P \|\nabla \cdot \boldsymbol{\sigma} + \mathbf{f}\|,$$

where C_P is the Poincaré constant of the compact embedding and $(\boldsymbol{\sigma}, q)$ is a reconstruction field pair. However, the drawback of this approach is that an expensive global minimization problem needs to be solved, e.g., for $(\boldsymbol{\sigma}, q)$ above. Another main reason to introduce extra field variables is that, for finite element methods, the basis functions are local and do not have a globally continuous derivative, in that $\Delta \mathbf{u}_S$ yields singular distributions, whose proper norm is H^{-1} -norm and cannot be evaluated by summing up element-wise L^2 -norms. In computation, it has to be replaced by $\nabla \cdot \boldsymbol{\sigma}$ where $\boldsymbol{\sigma} \in \mathbf{H}(\text{div})$, and is constructed to be closer to the true solution’s gradient $\nabla \mathbf{u}$ than $\nabla \mathbf{u}_S$. Meanwhile, for systems like NSE, the error estimation in Galerkin methods for NSE is further complicated by its saddle point nature from the divergence-free constraint, in that consistency has to be tweaked to ensure the Ladyzhenskaya-Babuška-Brezzi stability condition ([Girault & Raviart, 2012, Chapter III §1 Sec 1.2](#)). Recently, [Fanaskov et al. \(2024\)](#) considers NN as a function learner to represent solutions of linear convection-diffusion equations, yet still falls into the traditional error estimation framework in that an extra flux variable is learned by NN, and the error is of typical accuracy of NN function learners $\mathcal{O}(10^{-3})$. In contrast, in our study, the need of extra “stress” or “flux” variables to build the residual functional is circumvented as well.

Error correction for Bayesian inverse problems. The *a posteriori* error-correction approaches through sampling for the output of surrogate NNs in Bayesian inverse problems [Yan & Zhou \(2020\)](#). More recently, in [Cao et al. \(2023\)](#), the error equation in an inverse problem is approximated by solving for a Galerkin projection of a linearized error equation, while in this paper, we leverage the Gelfand triple directly to compute a nonlinear Galerkin projection directly by minimizing the spectral norm.

Sobolev norms in operator learning In the context of function learning problems, [Du et al. \(2024\)](#) applied an L^2 -spectral loss for PINN [Raissi et al. \(2019\)](#) and found superior results over mesh-weighted spatial losses. [Du et al. \(2024\)](#) realized the PDE residual as a functional on L^2 and exploited the Parseval identity on this space, which computes $\sup_{v \in L^2} \langle R, v \rangle / \|v\|_{L^2}$. We note this is not the natural space to measure the PDE residual according to the Hilbertian formulation of NSE.

The natural space to consider the residual functional is the dual space of $H^1(\mathbb{T}^2)$, where the solution to NSE belongs to $H^1(\mathbb{T}^2)$ at a given snapshot. The dual space of $H^1(\mathbb{T}^2)$ happens to be H^{-1} using the Gelfand triple. We compute $\sup_{v \in H^1} \langle R, v \rangle / \|v\|_{H^1}$. Using the Poisson equation $-\Delta u = f$ as an example, if $u \in H^1$, then the residual $f + \Delta u \in H^{-1}$. However, H^{-1} is a non-localizable norm in the spatial domain, and traditional wisdom (FEM or FVM) has to circumvent the direct evaluation of it by localizing the residual, which introduced inaccurate compromises in the residual-based error estimation.

A posteriori error estimation for flow problems There is a vast amount of literature for the *a posteriori* error estimation for the viscous flow ($1/\text{Re} > 0$) problems under the Hilbertian framework for traditional numerical methods. Our methods draw inspiration from these pioneers and try to address the drawbacks. The most popular *a posteriori* error estimator for stationary Stokes problem is from Verfürth (1989), and it is of residual-type by computing a mesh-weighted L^2 -norm elementwise, and the singular distribution Δu_S is represented by the magnitude of flux jumps across the facets in this mesh. In Bank & Welfert (1991), a more accurate *a posteriori* estimation technique is invented for Stokes problem in which a local problem is solved to represent the residual functional on a collection of neighboring elements. L^2 residual-type error estimation for stationary NSE is considered in Oden et al. (1994). To our best knowledge, no functional-type *a posteriori* error estimation has been applied to solve the transient NSE, due to its in-efficiency for traditional finite element or finite volume methods. As at every time step, a global nonlinear problem has to be solved if one ought to evaluate the functional accurately using finite element local basis functions, whose computational cost is an order of magnitude higher than implicit Euler methods.

Hybrid methods for turbulent flow predictions. There are quite a few approaches that combine NN-based learners with traditional time-marching solvers. Notable examples include: solution from Direct Numerical Simulation (DNS) corrected using an NN Um et al. (2020); interpolation to achieve spatial super-resolution using solutions produced by traditional numerical solvers marching on coarser grids Kochkov et al. (2021). orthogonal polynomial features (HiPPO from Gu et al. (2020)) to achieve temporal super-resolution in addition to the spatial one Sun et al. (2023). McGreivy & Hakim (2023) explores how to incorporate energy-preserving schemes such as upwinding into the ML solvers through traditional numerical methods. Qi & Sun (2024) identifies important frequency bands in FNO parametrization and improves FNO’s rolling-out performance in NSE benchmarks.

C DATA GENERATION

C.1 VORTICITY-STREAMFUNCTION FORMULATION

Here we derive the vorticity-streamfunction formulation for the convenience of the readers. Consider the standard NSE in 3D, \mathbf{u} ’s z -component is 0, and $\mathbf{u} = \mathbf{u}(t, x, y)$ has only planar dependence

$$\begin{cases} \frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\frac{1}{\rho} \nabla p + \nu \Delta \mathbf{u} + \mathbf{f}. \\ \nabla \cdot \mathbf{u} = 0 \end{cases} \quad (\text{C.1})$$

Taking $\nabla \times (\cdot)$ on both sides, and we assume that the solution is sufficiently regular that the spatial and temporal derivative can interchange, as well as curl and the Laplacian can interchange, one gets the following equation by letting $\boldsymbol{\omega} = \nabla \times \mathbf{u}$

$$\frac{\partial \boldsymbol{\omega}}{\partial t} + \nabla \times \nabla \left(\frac{\mathbf{u}^2}{2} \right) + \mathbf{u} \cdot \nabla \boldsymbol{\omega} - \boldsymbol{\omega} \cdot \nabla \mathbf{u} = -\nabla \times \left(\frac{1}{\rho} \nabla p \right) + \nu \Delta \boldsymbol{\omega} + \nabla \times \mathbf{f}. \quad (\text{C.2})$$

Next, one can have $\boldsymbol{\omega} \cdot \nabla \mathbf{u} = 0$, thus the equation becomes

$$\frac{\partial \boldsymbol{\omega}}{\partial t} + \mathbf{u} \cdot \nabla \boldsymbol{\omega} = \boldsymbol{\omega} \cdot \nabla \mathbf{u} + \nu \Delta \boldsymbol{\omega} + \nabla \times \mathbf{f}. \quad (\text{C.3})$$

We need to introduce a streamfunction ψ to recover the velocity $\mathbf{u} = \nabla \times (0, 0, \psi)$. The main simplification comes from the assumption of dependence on x - and y -variable only. Thus, in 2D this becomes $\mathbf{u} = \text{rot } \psi$ (a $\pi/2$ rotation of the 2D gradient $\nabla \psi$). Therefore,

$$\mathbf{u} = (u_1, u_2, 0) \implies u_1 = \partial_y \psi, u_2 = -\partial_x \psi$$




Figure 4: Ground truth streamlines for Taylor-Green vortex example.

and the vector vorticity can be equivalently represented by a scalar vorticity ω

$$\boldsymbol{\omega} = \nabla \times \mathbf{u} = (0, 0, \text{curl}(u_1, u_2)) = (0, 0, \partial_x u_2 - \partial_y u_1) =: (0, 0, \omega)$$

The equation becomes a nonlinear system for ω and ψ :

$$\begin{cases} \partial_t \omega + \mathbf{rot} \psi \cdot \nabla \omega - \nu \Delta \omega = \nabla \times \mathbf{f}, \\ \omega + \Delta \psi = 0. \end{cases} \quad (\text{C.4})$$

$$(\text{C.5})$$

C.2 TAYLOR-GREEN VORTEX

Taylor-Green vortex serves as one of the most well-known benchmarks for NSE numerical methods, as its flow type experiences from laminar, to transitional, and finally evolving into turbulent regime. We consider the trajectory before the breakdown phase. We opt to use a doubly periodic solution on $[0, 2\pi)^2$ such that the inflow/outflow occurs on the “boundary” (see Figure 4). The exact solution is given by:

$$\mathbf{u}(t, x, y) = e^{-2\kappa^2 \nu t} \begin{pmatrix} \sin(\kappa x) \cos(\kappa y) \\ -\cos(\kappa x) \sin(\kappa y) \end{pmatrix} \text{ or } e^{-2\kappa^2 \nu t} \begin{pmatrix} -\cos(\kappa x) \sin(\kappa y) \\ \sin(\kappa x) \cos(\kappa y) \end{pmatrix}$$

where $\nu = 1/\text{Re}$ is chosen as 10^{-3} . For a sample trajectory with $\kappa = 1$, please refer to Figure 4. We apply the pseudo-spectral method with RK2 time stepping for the convection term and Crank-Nicolson for the diffusion term. The dataset has 11 trajectories with κ ranging from 1 to 11, among which $\kappa = 11$ is chosen as the test trajectory.

C.3 ISOTROPIC TURBULENCE (KOLMOGOROV FLOW)

The example featured in the original FNO paper [Li et al. \(2020\)](#) can be viewed as a special case of the isotropic turbulence with a strongly regularized initial condition, no drag, and a force with low wavenumber. This is our example (I). The example (II) has weakly-regularized initial condition, small drag, and no external force. The initial condition spectra are chosen such that, after the warmup time, the spectra are the same with the Kolmogorov flow featured in the Jax-CFD paper [Kochkov et al. \(2021\)](#), in the sense of the inverse cascade in the Fourier domain (see Figure 8 and 9). The training data used in both (I) and (II) are generated using a pseudo-spectral method with a de-aliasing filter using an RK4 marching for the convection term and implicit Euler for the diffusion term. The time steps for both are $\Delta t = 10^{-3}$ on a 256^2 grid, then downsampled to a 64^2 grid. The warmup time is 4.5. The δt for the ST-FNO training and evaluations are $\delta t = 55\Delta t$. The testing data are generated on a 512^2 grid with a 5×10^{-4} time step.




Figure 5: Contours plots of pointwise values of residuals for Example (I). (a): the residual of the ground truth; (b): residual of ST-FNO, where the time derivative in the residual is using the ground truth’s; (c): the residual after fine-tuning for 10 ADAM iterations where the time derivative is computed using an extra-fine-step numerical solver.




Figure 6: Pointwise values of residual of ST-FNO in the frequency domain in \ln scale of Example (I). (a): the residual of ST-FNO before correction, time derivative computed using the ground truth; (b): the residual of ST-FNO after fine-tuning for 10 ADAM iterations where the time derivative is computed using an extra-fine-step numerical solver.

D EXPERIMENTS

D.1 TRAINING AND EVALUATION

The training uses `torch.optim.OneCycleLR` Smith & Topin (2019) learning rate strategy with a 20% warm-up phase. AdamW is the optimizer with no extra regularization. The learning rate starts and ends with $10^{-3} \cdot lr_{\max}$. The $lr_{\max} = 10^{-3}$. The result demonstrated is obtained from fixing the random number generator seed. All models are trained on a single RTX A4500 or RTX A6000. The codes to replicate the experiments are open-source and publicly available.³

Despite using the negative Sobolev norm is quite efficient in fine-tuning, using it in training not be the most efficient in minimizing the L^2 -norm due to the optimization being non-convex, we observe that all norms are equivalent “bad” due to nonlinearity of NOs.

D.2 MODELS

Helmholtz layer for the velocity-pressure formulation. For the V-P formulation in a simplified connected convex or periodic domain, it is known that an exact divergence-free subspace $\mathcal{W} \subset \mathcal{V}$ for velocity means that the pressure field is not needed. The reason is that in the weak formulation, the pressure is a Lagrange multiplier to impose the divergence-free condition (Girault & Raviart, 2012, Chapter III 1 Section 1.1). Inspired by the postprocessing to eliminate ∇p (Ku et al., 1987, eq. (15)) together with the neural Clifford layers Brandstetter et al. (2023), we add a Helmholtz layer S after each application of $\sigma_j \circ (W_j + K_j)$. S performs a discrete Helmholtz decomposition (Girault & Raviart, 2012, Chapter 1 §1 Section 3.1) in the frequency domain to project the latent fields to be solenoidal. For details on the implementation, please refer to Appendix D.2.

³Please download the code in the anonymous repository and follow the instructions in `README.md`.




Figure 7: A simple illustration of necessity of padding using a 1D non-periodic data. FNO3d cannot handle temporal non-periodic data. The free temporal super-resolution by FFT/iFFT has Runge-like phenomena near the end points without padding. Padding in the ST-FNO3d modification resolves this problem. Since data themselves are intrinsically lower-dimensional, we opt not to pad the latent representations, but rather apply the padding only in the last spectral convolution layer to get the output.

Table 6: The detailed comparison of FNO3d and its spatiotemporal modification. Layer: # of spectral convolution layers; SFNO has an extra layer of spectral convolution in a single channel with skip connection. Modes: (τ_{\max}, k_{\max}) . Pre-norm: whether a pointwise Gaussian normalizer is applied for the input data. Eval FLOPs: Giga FLOPs for one evaluation instance. FT: finetuning GFLOPs profiles per ADAM iterations for a 256^2 grid for a single instance. A `torch.cfloat` type parameter entry counts as two parameters.

	Architectures					GFLOPs		# params
	layers	channel/width	modes	activation	pre-norm	Eval	FT	
FNO3d Example 2 (I)	4	20	(5, 8)	GELU	Y	13.3	N/A	9.03m
ST-FNO3d Example 2 (I)	$4 + \frac{1}{20}$	20	(5, 8)	GELU	N	27.0	3.28	9.02m
FNO3d Example 2 (II)	4	10	(5, 32)	GELU	Y	28.7	N/A	16.38m
ST-FNO3d Example 2 (II)	$4 + \frac{1}{10}$	10	(5, 32)	GELU	N	42.5	3.4	16.42m

Difference with Fourier Neural Operator 3d In view of (2.5), the FFTs in FNO3d transforms are continuous integrals in the spatial dimensions yet a discrete sum in the temporal dimension. Despite that no explicit restriction imposed on the output time steps, the d_{out} cannot be trivially changed as the data are prepared by applying a pointwise Gaussian normalizer that depends on d_{out} . Denote the lifting operator (channel expansion) in FNO3d by $P : \{\mathbf{a}_h\} \oplus \{\mathbf{p}_h\} \rightarrow \mathbb{R}^{d_v \times d_{\text{out}} \times n \times n}$. Before the application of P in FNO3d, \mathbf{a}_h is artificially repeated d_{out} times, then P is in a space of linear operators $\simeq \mathbb{R}^{(d_{\text{in}}+3) \times d_v}$. This observation makes it independent of the spatial discretization size $n \times n$ by dependent on d_{in} . In FNO3d, the channel reduction is a pointwise nonlinear universal approximator, yet in ST-FNO3d, this is linear, which eventually facilitate the convexity of the fine-tuning optimization problem.

D.3 FINE-TUNING

In implementations, we choose $\text{Iter}_{\max} = 100$ and the ADAM optimizer with learning rate 10^{-1} in Algorithm 1, i.e., an ADAM is simply run for 100 iterations to update parameters that count only a fraction of a spectral layer since there is only 1 channel of weights.

To train or to stop? Investigating the frequency domain signatures In Stuart’s paper on the convergence of linear operator learning de Hoop et al. (2023), the evaluation error scales with number of samples, in the few-data regime, the overfitting comes fast. Part of the reason is that the operator learning problems are in the “small” data regime.

Denote $\omega(t, \mathbf{x}) := \nabla \times \mathbf{u}$, and $\hat{\omega} := \mathcal{F}\omega$ where \mathcal{F} is applied only in space, then we can define enstrophy spectra as follows:

$$\mathcal{E}(t, k) = \sum_{k-\delta k \leq |\mathbf{k}| \leq k+\delta k} |\hat{\omega}(t, \mathbf{k})|^2, \text{ and } \sum_k \mathcal{E}(t, k) := \int_{\Omega} |\omega|^2 d\mathbf{x}$$

In Example (II), the enstrophy spectrum decays as $\mathcal{O}(k^{-3})$, which is slightly faster than kinetic energy. This is usually referred to as the direct cascade, as opposed to the inverse cascade $\mathcal{O}(k^{-5/3})$ discovered by Kolmogorov. For SFNO, it learns the frequency signature of the data after a single epoch. See Figure 8 and Figure 9.




Figure 8: Enstrophy spectrum density comparison for Example (II) to illustrate the “convergence” of the SFNO training. The SFNO evaluation is on a 256×256 grid for a fixed randomly chosen sample, and the training is on a 64×64 grid starting from 10 different seeds at different time steps. The error bars are plotted with $+/-$ 10 times the standard deviation from the mean to boost the visibility of the convergence. (a): the comparison after 1 epoch, the average relative L^2 difference with the ground truth is 2.150×10^{-1} ; (b): the comparison after 10 epochs, the average relative L^2 difference with the ground truth is 8.051×10^{-2} .

D.4 ADDITIONAL EXPERIMENTS

In this subsection, we perform the following additional experiments regarding the possibility of fine-tuning ST-FNO3d for “out-of-distribution” data. All ST-FNO3d models are trained for 15 epochs (5 more than the model in Table 3) with the data generated using the McWilliams’s isotropic turbulence data from Example (II) with $\nu = 10^{-3}$. The initial energy distribution satisfies:

$$|\hat{\psi}(k)|^2 \sim k^{-1}(\tau^2 + (k/k_0)^4)^{-1}. \quad (\text{D.1})$$

The parameter k_0 in Example (II) roughly means that the initial energy is concentrated at $|\mathbf{k}| \approx k_0$. In Table 3 k_0 is chosen as 4 according to McWilliams (1984). For all examples, the same “burn-in” time is used ($t = 4.5$), the training output time steps are 10 and evaluation output time steps are 40. Please refer to Figure 10 for sample trajectories and Figure 11 for the energy cascade after the “burn-in” time period.

- (Ex1) The initial energy is concentrated at $k_0 = 2$. A sample trajectory can be found in Figure 10(a).
- (Ex2) Same with above, $k_0 = 8$. A sample trajectory can be found in Figure 10(c).
- (Ex3) $k_0 = 4$, $\nu = 1/Re = 1/5000$. In order to resolve the small viscosity, the mesh size becomes 512×512 .
- (Ex4) $k_0 = 8$, $\nu = 1/Re = 1/10000$, mesh size is 1024×1024 for the same reason as above.




Figure 9: Enstrophy spectrum density comparison for Example (II) for two randomly selected trajectories for a given trained ST-FNO. Training and evaluation setups are identical to those in Figure 8. (a), (b), (c): the comparison of the spectra of the ST-FNO evaluation and the ground truth after 10 epochs of training at $t_{\ell+4}$, $t_{\ell+7}$, and $t_{\ell+10}$ of the first trajectory; (d), (e), (f): the comparison for the second trajectory at the same time steps.

The evaluation and fine-tuning result can be found in Table 7. We found that the new paradigm works reasonably well for first three example, yet fails to converge for the last example. For an example plot of the prediction comparison please refer to Figure 12. The correlation graph between the ST-FNO3d’s predictions and the ground truth can be found in Figure 13, in which the FNO(2+1)d is adapted from the original FNO2d code. This roll-out-based model predicts the next snapshot based on 10 given snapshots as input, and the predicted snapshot is concatenated to the current input as the latest snapshot to form the input for the next prediction.

Table 7: Evaluation metrics of the McWilliams isotropic turbulence example for “out-of-distribution” data. For details please refer to Section D.4 $\varepsilon := \omega_S - \omega_N$. $r := R(\omega_N)$. $\mathcal{Y} := H^{-1}(\mathbb{T}^2)$. All error norms are evaluated at the final time step. The training is under the \tilde{L}^2 -norm and the fine-tuning is under the H^{-1} -norm. Symbol “-” means the errors are the same as the entry on the row above.

	In-distribution evaluation		Out-of-distribution evaluation		After fine-tuning	
	$\ \varepsilon\ _{L^2}$	$\ R\ _{-1,n}$	$\ \varepsilon\ _{L^2}$	$\ R\ _{-1,n}$	$\ \varepsilon\ _{L^2}$	$\ R\ _{-1,n}$
(Ex1) $k_0 = 2$	5.39×10^{-2}	1.66×10^{-2}	8.52×10^{-2}	4.25×10^{-2}	2.73×10^{-2}	5.05×10^{-5}
(Ex2) $k_0 = 8$	-	-	1.46×10^{-1}	2.31×10^{-2}	2.06×10^{-2}	5.85×10^{-6}
(Ex3) $Re = 5 \cdot 10^3$	-	-	2.94×10^{-1}	2.73×10^{-2}	2.50×10^{-2}	2.96×10^{-4}
(Ex4) $Re = 10^4$	-	-	4.92×10^{-1}	1.76×10^{-1}	2.355×10^6	2.873×10^4




Figure 10: Examples trajectories of the evaluation datasets for additional experiments for Example (II). (a): $k_0 = 2$; (b): $k_0 = 4$. (c): $k_0 = 8$.




Figure 11: Energy spectrum densities datasets for additional experiments for Example (II). (a): $k_0 = 2$; (b): $k_0 = 4$. (c): $k_0 = 8$.

E ASSUMPTIONS AND PROOFS

Assumption E.1 (Assumptions for Theorems 3.1, 3.2, E.9). *The following notions and assumptions are adopted throughout the proof of the three theorems involved, all of which are common in the literature for NSE. While some of them can be proved, we opt for list them here to facilitate the presentation.*

(E₁) *The compact embeddings for Gelfand triple $\mathcal{V} \Subset \mathcal{H} \Subset \mathcal{V}'$ hold, where $\mathcal{V} = \mathbf{H}^1(\mathbb{T}^2)$, and $\mathcal{H} = \mathbf{L}^2(\mathbb{T}^2)$.*

(E₂) *The time interval \mathcal{T} in consideration is fixed in that we consider the “refining” problem for a fixed time interval but with variable time steps.*




Figure 12: The ground truth versus the ST-FNO3d evaluation/fine-tuning results for the $Re = 5000$ additional experiment for Example (II). (a): the ground truth at $t = 5$ (prediction step 20) generated by IMEX RK4; (b): ST-FNO3d prediction. (c): ST-FNO3d fine-tuned prediction.




Figure 13: Comparison of the correlation with the ground truth (on 1024×1024 grid) generated by IMEX RK4 solver.

(E₃) The initial condition $\mathbf{u}_0 \in \mathcal{V}$, and $\mathbf{f} \in L^2(\mathcal{T}; \mathcal{V}')$.

Lemma E.2 (Skew-symmetry of the trilinear term). For $\mathbf{z}, \mathbf{u}, \mathbf{v} \in \mathbf{H}^1(\mathbb{T}^2)$, where \mathbb{T}^2 denotes $\Omega := (0, 1)^2$ equipped with component-wise PBC. Denote

$$c(\mathbf{z}, \mathbf{u}, \mathbf{v}) := ((\mathbf{z} \cdot \nabla) \mathbf{u}, \mathbf{v}) = \int_{\Omega} ((\mathbf{z} \cdot \nabla) \mathbf{u}) \cdot \mathbf{v} \, d\mathbf{x}.$$

If $\nabla \cdot \mathbf{z} = 0$, then

$$c(\mathbf{z}, \mathbf{u}, \mathbf{v}) = -c(\mathbf{z}, \mathbf{v}, \mathbf{u}), \quad (\text{E.1})$$

and specifically $c(\mathbf{z}, \mathbf{v}, \mathbf{v}) = 0$.

Proof. We note that this result is common in NSE literature, see e.g., (Temam, 1995, Part I Sec 2.3), for homogeneous boundary or whole space. In this case, it suffices to verify that the boundary term vanishes. Here we still provide a short argument for the convenience of readers. First by the product rule, and a vector calculus identity of $\nabla \cdot (\mathbf{v} \cdot \mathbf{u})$ (see e.g., (Balanis, 2012, Appendix II.3.2))

$$\nabla \cdot ((\mathbf{v} \cdot \mathbf{u}) \mathbf{z}) = ((\mathbf{z} \cdot \nabla) \mathbf{u}) \cdot \mathbf{v} + ((\mathbf{z} \cdot \nabla) \mathbf{v}) \cdot \mathbf{u} + (\mathbf{v} \cdot \mathbf{u}) \nabla \cdot \mathbf{z}. \quad (\text{E.2})$$

By Gauss divergence theorem, we then obtain:

$$\int_{\Omega} \nabla \cdot ((\mathbf{v} \cdot \mathbf{u}) \mathbf{z}) \, d\mathbf{x} = \int_{\partial\Omega} (\mathbf{v} \cdot \mathbf{u}) \mathbf{z} \cdot \mathbf{n} \, ds = \sum_{e_i \subset \partial\Omega, 1 \leq i \leq 4} \int_{\partial\Omega} (\mathbf{v} \cdot \mathbf{u}) \mathbf{z} \cdot \mathbf{n}_i \, ds,$$

with \mathbf{n} is outer normal with respect to $\partial\Omega$, and \mathbf{n}_i denotes that with respect to edge e_i . Now on $e_1 := \{x = 0\} \times \{y \in (0, 1)\}$, $\mathbf{n}_1 = (-1, 0)^\top$, $\mathbf{w}|_{e_1} = \mathbf{w}|_{e_2}$ by PBC where $e_2 := \{x = 1\} \times \{y \in (0, 1)\}$ and $\mathbf{w} \in \{\mathbf{z}, \mathbf{u}, \mathbf{v}\}$. The integrals on e_1 and e_2 cancel with one another since $\mathbf{n}_2 = -\mathbf{n}_1$. Furthermore, if $\nabla \cdot \mathbf{z} = 0$, integrating (E.2) on Ω yields the desired result. \square

Lemma E.3 (Poincaré inequality). $\mathcal{V}/\mathbb{R} \in \mathcal{H}/\mathbb{R}$ is compact and the following Poincaré inequality holds with constant 1: for any $v \in \mathcal{V}/\mathbb{R}$

$$\|v\| \leq |v|_1 \quad (\text{E.3})$$

Proof. A common textbook proof exploits the equivalence of sequential compactness and compactness, once the norm topology is introduced, we prove this in a very intuitive way under our special setting: by definition (3.7), if $v \in \mathcal{V}/\mathbb{R}$ then $\hat{v}(0, 0) = 0$

$$\|v\|_0^2 = \sum_{\mathbf{k} \in \mathbb{Z}^2 \setminus \{\mathbf{0}\}} |\hat{v}(\mathbf{k})|^2 \leq \sum_{\mathbf{k} \in \mathbb{Z}^2 \setminus \{\mathbf{0}\}} |\mathbf{k}|^2 |\hat{v}(\mathbf{k})|^2 = |v|_1^2.$$

\square

E.1 PROOF OF THEOREM 3.1

To prove Theorem 3.1, we need Lemmas E.4, E.5, and E.6, all of which can be found in classical textbook on analysis and approximations of NSE, such as Temam (1995), or Girault & Raviart (2012) for homogeneous Dirichlet boundary condition for the velocity field or slip boundary condition for the (pseudo-)stress tensor. We shall complement some proofs with simple adaptations to PBC whenever necessary for the convenience of readers.

Before presenting any lemmas, we need the following definitions, a weighted \mathbf{H}^1 -norm on $\mathcal{V} = \mathbf{H}^1(\mathbb{T}^2)$ for $\alpha > 0$ on Ω is defined as

$$\|\mathbf{v}\|_{\alpha,1} := (\|\mathbf{v}\|_{L^2}^2 + \|\alpha \nabla \mathbf{v}\|_{L^2}^2)^{1/2}, \text{ and } |\mathbf{v}|_{\alpha,1} := \|\alpha \nabla \mathbf{v}\|_{L^2}.$$

For $\mathbf{f} \in \mathcal{V}'$, define a weighted dual norm as follows:

$$\|\mathbf{f}\|_{\alpha^{-1},-1} := \sup_{\mathbf{v} \in \mathcal{V}} \frac{\langle \mathbf{f}, \mathbf{v} \rangle}{\|\mathbf{v}\|_{\alpha,1}} \quad (\text{E.4})$$

Lemma E.4 (Continuity and embedding results for the convection term). *For trilinear convection term*

$$((\mathbf{u} \cdot \nabla) \mathbf{v}, \mathbf{w}) \leq \|\mathbf{u}\|_{L^4} |\mathbf{v}|_1 \|\mathbf{v}\|_{L^4}. \quad \text{and} \quad ((\mathbf{u} \cdot \nabla) \mathbf{v}, \mathbf{w}) \leq |\mathbf{u}|_1 |\mathbf{v}|_1 |\mathbf{w}|_1 \quad (\text{E.5})$$

Proof. The first result is obtained by the Hölder inequality, and the second holds with suite Sobolev embedding results (Temam, 1995, Part I, §2.3, Lemma 2.1). \square

Lemma E.5 (Energy stability of NSE in a bounded domain).

$$\|\mathbf{u}\|_{L^\infty(\mathcal{T}; \mathcal{H})}^2 \leq \|\mathbf{u}_0\|_{\mathcal{H}}^2 + \|\mathbf{f}\|_{L^2(\mathcal{T}; \mathcal{V}')}^2. \quad (\text{E.6})$$

Proof. (E.17) is a classical result, see e.g., (Temam, 1995, Section 3.1). \square

Lemma E.6 (Fréchet derivative of the convection term). *Given $\mathbf{v}, \mathbf{u} \in \mathcal{V}$, the linearization of the difference of the convection term reads*

$$\mathbf{v} \cdot \nabla \mathbf{v} = \mathbf{u} \cdot \nabla \mathbf{u} + D(\mathbf{u})(\mathbf{v} - \mathbf{u}) + \mathbf{r} \text{ where } D(\mathbf{u})\xi := (\xi \cdot \nabla) \mathbf{u} + (\mathbf{u} \cdot \nabla) \xi, \quad (\text{E.7})$$

is the Fréchet derivative and $\|\mathbf{r}\| \lesssim \|\mathbf{u} - \mathbf{v}\| \|\mathbf{u} - \mathbf{v}\|_1$.

Proof. This result is normally used without proof in linearizing NSE, see e.g., (Girault & Raviart, 2012, Chapter 4 eq. (6.5.2)). Define $F(\mathbf{u}, \mathbf{G}) := \mathbf{u}^\top \mathbf{G}$, then expanding F at \mathbf{v} and $\mathbf{G} = \nabla \mathbf{u}$ yields the desired result. \square

Theorem 2.1 Part I (Functional-type a posteriori error estimate is efficient (rigorous version)). Consider the Gelfand triple $\mathcal{V} \Subset \mathcal{H} \Subset \mathcal{V}'$, and the weak solution $\mathbf{u} \in L^2(\mathcal{T}; \mathcal{V}) \cap L^\infty(\mathcal{T}; \mathcal{H})$ to (2.2) be sufficiently regular, then the dual norm of the residual is efficient to estimate the error:

$$\|\mathcal{R}(\mathbf{u}_\mathcal{N})\|_{L^2(\mathcal{T}; \mathcal{V}')}^2 \lesssim \|\partial_t(\mathbf{u} - \mathbf{u}_\mathcal{N})\|_{L^2(\mathcal{T}; \mathcal{V}')}^2 + \|\mathbf{u} - \mathbf{u}_\mathcal{N}\|_{L^2(\mathcal{T}; \mathcal{V})}^2. \quad (\text{E.8})$$

The constant depend on the regularity of the true solution \mathbf{u} .

Proof. The proof of the lower bound (efficiency) (E.8) follows the skeleton laid out in (Verfürth, 2003, Lemma 4.1) for diffusion and temporal derivative terms, while not needing to extend approximation using an affine linear extension operator as Verfürth (2003) does. The treatment of the nonlinear convection term follows (Fischer, 2015, Lemma 7) for the steady-state NSE (no temporal derivatives), barring the technicality of the divergence-free condition. Consider at $t \in \mathcal{T}$, for any test function $\mathbf{v} \in \mathbf{H}^1(\mathbb{T}^2)$ and $\nabla \cdot \mathbf{v} = 0$, integrating \mathbf{v} against $R(\mathbf{u}_\mathcal{N})$ in (3.2) yields

$$\begin{aligned} \langle R(\mathbf{u}_\mathcal{N}), \mathbf{v} \rangle &= (\mathbf{f} - \partial_t \mathbf{u}_\mathcal{N} - (\mathbf{u}_\mathcal{N} \cdot \nabla) \mathbf{u}_\mathcal{N} + \nu \Delta \mathbf{u}_\mathcal{N}, \mathbf{v}) \\ &= (\partial_t(\mathbf{u} - \mathbf{u}_\mathcal{N}), \mathbf{v}) - (\nu \Delta(\mathbf{u} - \mathbf{u}_\mathcal{N}), \mathbf{v}) \\ &\quad + ((\mathbf{u} \cdot \nabla) \mathbf{u} - (\mathbf{u}_\mathcal{N} \cdot \nabla) \mathbf{u}_\mathcal{N}, \mathbf{v}). \end{aligned}$$

Integrating by parts for the diffusion term, using the same argument for PBC as in Lemma E.2, and inserting $(\mathbf{u} \cdot \nabla) \mathbf{u}_\mathcal{N}$ yield

$$\begin{aligned} \langle R(\mathbf{u}_\mathcal{N}), \mathbf{v} \rangle &= \underbrace{(\partial_t(\mathbf{u} - \mathbf{u}_\mathcal{N}), \mathbf{v})}_{\text{(I)}} + \underbrace{(\nu \nabla(\mathbf{u} - \mathbf{u}_\mathcal{N}), \nabla \mathbf{v})}_{\text{(II)}} \\ &\quad + \underbrace{((\mathbf{u} \cdot \nabla)(\mathbf{u} - \mathbf{u}_\mathcal{N}), \mathbf{v})}_{\text{(III)}} + \underbrace{(((\mathbf{u} - \mathbf{u}_\mathcal{N}) \cdot \nabla) \mathbf{u}_\mathcal{N}, \mathbf{v})}_{\text{(IV)}}. \end{aligned} \quad (\text{E.9})$$

Applying the definition of a weighted dual norm (E.4) on (I) we have

$$\text{(I)} \leq \|\partial_t(\mathbf{u} - \mathbf{u}_\mathcal{N})(t, \cdot)\|_{\nu^{-1/2}, -1} \|\mathbf{v}\|_{\nu^{1/2}, 1}.$$

For (II) we simply have

$$\text{(II)} \leq \|\nu^{1/2} \nabla(\mathbf{u} - \mathbf{u}_\mathcal{N})\|_{L^2} \|\nu^{1/2} \nabla \mathbf{v}\|_{L^2}.$$

For (III), applying Lemma E.4 and the stability estimate in Lemma E.5 for \mathbf{u} :

$$\text{(III)} \leq C \|\mathbf{u}\|_{\nu^{-1}, 1} \|\mathbf{u} - \mathbf{u}_\mathcal{N}\|_{\nu^{1/2}, 1} \|\mathbf{v}\|_{\nu^{1/2}, 1} \leq \nu^{-1} C_1(t, \mathbf{u}_0) \|\mathbf{u} - \mathbf{u}_\mathcal{N}\|_{\nu^{1/2}, 1} \|\mathbf{v}\|_{\nu^{1/2}, 1}.$$

For (IV), applying Lemma E.4 and the stability estimate in Lemma E.5 for $\mathbf{u}_\mathcal{N}$:

$$\text{(IV)} \leq C \|\mathbf{u} - \mathbf{u}_\mathcal{N}\|_{\nu^{1/2}, 1} \|\mathbf{u}_\mathcal{N}\|_{\nu^{-1}, 1} \|\mathbf{v}\|_{\nu^{1/2}, 1} \leq \nu^{-1} C_2(t, \mathbf{u}_0) \|\mathbf{u} - \mathbf{u}_\mathcal{N}\|_{\nu^{1/2}, 1} \|\mathbf{v}\|_{\nu^{1/2}, 1}.$$

Applying the Cauchy-Schwarz inequality, one has

$$|\langle R(\mathbf{u}_\mathcal{N}), \mathbf{v} \rangle|^2 \leq C \left(\|\partial_t(\mathbf{u} - \mathbf{u}_\mathcal{N})(t, \cdot)\|_{\nu^{-1/2}, -1}^2 + \|\mathbf{u} - \mathbf{u}_\mathcal{N}\|_{\nu^{1/2}, 1}^2 \right) \|\mathbf{v}\|_{\nu^{1/2}, 1}^2.$$

Finally, equipping \mathcal{V}' with the weighted norm (E.4), and by the definition of $L^2(\mathcal{T}; \mathcal{V}')$, we have

$$\begin{aligned} \|R(\mathbf{u}_\mathcal{N})\|_{L^2(\mathcal{T}; \mathcal{V}')}^2 &= \int_{\mathcal{T}} \sup_{\mathbf{v} \in \mathcal{V}'} \frac{|\langle R(\mathbf{u}_\mathcal{N}), \mathbf{v} \rangle|^2}{\|\mathbf{v}\|_{\nu^{1/2}, 1}^2} dt \\ &\leq C \int_{\mathcal{T}} \left(\|\partial_t(\mathbf{u} - \mathbf{u}_\mathcal{N})(t, \cdot)\|_{\nu^{-1/2}, -1}^2 + \|\mathbf{u} - \mathbf{u}_\mathcal{N}\|_{\nu^{1/2}, 1}^2 \right) dt, \end{aligned}$$

where $C = C(\mathbf{u}_0, \nu, \mathcal{T})$. \square

Theorem 2.1 Part II (Functional-type a posteriori error estimate is reliable (rigorous version)). Consider the Gelfand triple $\mathcal{V} \Subset \mathcal{H} \Subset \mathcal{V}'$, and the weak solution is $\mathbf{u} \in L^2(\mathcal{T}; \mathcal{V}) \cap L^\infty(\mathcal{T}; \mathcal{H})$ to (2.2). We assume that

(2.1.1) $\mathbf{u}_\mathcal{N}$ and \mathbf{u} are sufficiently close in the sense that: there exists $\gamma \in [0, C)$ for a fixed $C \in \mathbb{R}^+$, for \mathbf{J} defined in Lemma E.6

$$\|\mathbf{J}(\mathbf{u}, \mathbf{u}_\mathcal{N})\|_{\mathcal{V}'} \leq \gamma \|\nabla(\mathbf{u} - \mathbf{u}_\mathcal{N})\|. \quad (\text{E.10})$$

(2.1.2) \mathbf{u} satisfies the Gårding inequality (a weaker coercivity): for any $\mathbf{v}(t, \cdot) \in \mathcal{V}$, define

$$B(\mathbf{v}, \mathbf{w}; \mathbf{u}) := (\partial_t \mathbf{v}, \mathbf{w}) + \nu(\nabla \mathbf{v}, \nabla \mathbf{w}) + ((\mathbf{u} \cdot \nabla) \mathbf{v}, \mathbf{w}),$$

and for \mathbf{v} in a neighborhood such that Assumption (2.1.1) holds, there exists $\alpha, \beta > 0$ such that $\alpha - \beta \geq \nu$,

$$B(\mathbf{v}, \mathbf{v}; \mathbf{u}) + \beta \|\mathbf{v}\|^2 \geq \frac{d}{dt} \|\mathbf{v}\|^2 + \alpha \|\nabla \mathbf{v}\|^2. \quad (\text{E.11})$$

The dual norm of the residual is then reliable to serve as an error measure in the following sense: denote $\mathcal{T}_m := (t_m, t_{m+1}]$

$$\|\mathbf{u} - \mathbf{u}_{\mathcal{N}}\|_{L^\infty(\mathcal{T}_m; \mathcal{H})}^2 + \|\mathbf{u} - \mathbf{u}_{\mathcal{N}}\|_{L^2(\mathcal{T}_m; \mathcal{V})}^2 \leq \|(\mathbf{u} - \mathbf{u}_{\mathcal{N}})(t_m, \cdot)\|_{\mathcal{V}}^2 + C \int_{\mathcal{T}_m} \|R(\mathbf{u}_{\mathcal{N}})(t, \cdot)\|_{\mathcal{V}}^2 dt. \quad (\text{E.12})$$

The constant $C = C(\nu)$.

Proof. The proof of part II of Theorem 3.1 combines the meta-framework of Verfürth (2003) for time-dependent problems and (Verfürth, 1994, Section 8) for the stationary NSE. Note that thanks for the construction of divergence-free $\mathbf{u}_{\mathcal{N}}$, the technicality of applying the Ladyzhenskaya-Babuška-Brezzi inf-sup condition is avoided. In the meantime, no interpolations are needed to convert the functional norm for the negative Sobolev spaces to an L^2 representation.

To prove the upper bound (reliability) (3.5), we simply choose a time step $t = t_m$, and let $\mathbf{v} = (\mathbf{u} - \mathbf{u}_{\mathcal{N}})(t, \cdot)$ on (t_m, t_{m+1}) . Note that the discretized approximation's evaluation at t , which may not be the temporal grids t_m or t_{m+1} are naturally defined using the basis in (3.1). Using the error representation in (E.9), we have the (I) and (II) terms in $\langle R(\mathbf{u}_{\mathcal{N}}), \mathbf{u} - \mathbf{u}_{\mathcal{N}} \rangle$ are

$$\begin{aligned} \text{(I)} + \text{(II)} &= (\partial_t(\mathbf{u} - \mathbf{u}_{\mathcal{N}}), \mathbf{u} - \mathbf{u}_{\mathcal{N}}) + (\nu \nabla(\mathbf{u} - \mathbf{u}_{\mathcal{N}}), \nabla(\mathbf{u} - \mathbf{u}_{\mathcal{N}})) \\ &= \frac{1}{2} \frac{d}{dt} \|\mathbf{u} - \mathbf{u}_{\mathcal{N}}\|_{L^2}^2 + \nu \|\nabla(\mathbf{u} - \mathbf{u}_{\mathcal{N}})\|_{L^2}^2 \end{aligned} \quad (\text{E.13})$$

For the (III) and (IV) terms we have

$$\text{(III)} + \text{(IV)} = \overbrace{((\mathbf{u} \cdot \nabla)(\mathbf{u} - \mathbf{u}_{\mathcal{N}}), \mathbf{u} - \mathbf{u}_{\mathcal{N}}))}^{=0 \text{ by Lemma E.2}} + \overbrace{(((\mathbf{u} - \mathbf{u}_{\mathcal{N}}) \cdot \nabla) \mathbf{u}_{\mathcal{N}}, \mathbf{u} - \mathbf{u}_{\mathcal{N}}))}^{=:(*)}. \quad (\text{E.14})$$

Note by Lemma E.2, $(*)$ is also:

$$(*) = (((\mathbf{u} - \mathbf{u}_{\mathcal{N}}) \cdot \nabla) \mathbf{u}, \mathbf{u} - \mathbf{u}_{\mathcal{N}}), \quad (\text{E.15})$$

since $c(\mathbf{z}, \mathbf{z}, \mathbf{z}) = 0$ for $\mathbf{z} = \mathbf{u} - \mathbf{u}_{\mathcal{N}}$. Consequently, by a vector calculus identity, we have

$$\int_{\Omega} (\mathbf{z} \cdot \nabla) \mathbf{u} \cdot \mathbf{z} \, d\mathbf{x} = \int_{\Omega} \nabla \mathbf{u} : (\mathbf{z} \otimes \mathbf{z}) \, d\mathbf{x}.$$

Thus, we reach the following error equation:

$$\langle R(\mathbf{u}_{\mathcal{N}}), \mathbf{u} - \mathbf{u}_{\mathcal{N}} \rangle = \frac{1}{2} \frac{d}{dt} \|\mathbf{u} - \mathbf{u}_{\mathcal{N}}\|_{L^2}^2 + \nu \|\nabla(\mathbf{u} - \mathbf{u}_{\mathcal{N}})\|_{L^2}^2 + \int_{\Omega} \nabla \mathbf{u} : ((\mathbf{u} - \mathbf{u}_{\mathcal{N}}) \otimes (\mathbf{u} - \mathbf{u}_{\mathcal{N}})) \, d\mathbf{x}. \quad (\text{E.16})$$

Now, by Assumption (2.1.2), Poincaré inequality from Lemma E.3,

$$\begin{aligned} \frac{1}{2} \frac{d}{dt} \|\mathbf{u} - \mathbf{u}_{\mathcal{N}}\|^2 + \nu \|\nabla(\mathbf{u} - \mathbf{u}_{\mathcal{N}})\|^2 &\leq \frac{d}{dt} \|\mathbf{u} - \mathbf{u}_{\mathcal{N}}\|^2 + (\alpha - \beta) \|\nabla(\mathbf{u} - \mathbf{u}_{\mathcal{N}})\|^2 \\ &\leq \frac{d}{dt} \|\mathbf{u} - \mathbf{u}_{\mathcal{N}}\|^2 + \alpha \|\nabla(\mathbf{u} - \mathbf{u}_{\mathcal{N}})\|^2 - \beta \|\mathbf{u} - \mathbf{u}_{\mathcal{N}}\|^2 \\ &\leq B(\mathbf{u} - \mathbf{u}_{\mathcal{N}}, \mathbf{u} - \mathbf{u}_{\mathcal{N}}; \mathbf{u}). \end{aligned}$$

By (E.13), (E.14), and (E.15),

$$\begin{aligned} B(\mathbf{u} - \mathbf{u}_{\mathcal{N}}, \mathbf{u} - \mathbf{u}_{\mathcal{N}}; \mathbf{u}) &= \text{(I)} + \text{(II)} + \text{(III)} + \text{(IV)} = \langle R(\mathbf{u}_{\mathcal{N}}), \mathbf{u} - \mathbf{u}_{\mathcal{N}} \rangle \\ &\leq \|R(\mathbf{u}_{\mathcal{N}})\|_{\nu^{-1/2}, -1} \|\mathbf{u} - \mathbf{u}_{\mathcal{N}}\|_{\nu^{1/2}, 1} \\ &\leq \frac{1}{2} \|R(\mathbf{u}_{\mathcal{N}})\|_{\nu^{-1/2}, -1}^2 + \frac{1}{2} \|\mathbf{u} - \mathbf{u}_{\mathcal{N}}\|_{\nu^{1/2}, 1}^2. \end{aligned}$$

As a result, we have

$$\frac{d}{dt} \|\mathbf{u} - \mathbf{u}_{\mathcal{N}}\|^2 + \nu \|\nabla(\mathbf{u} - \mathbf{u}_{\mathcal{N}})\|^2 \leq \|R(\mathbf{u}_{\mathcal{N}})\|_{\nu^{-1/2}, -1}^2. \quad (\text{E.17})$$

Integrating from t_m to any $t \in (t_m, t_{m+1}]$ yields

$$\|(\mathbf{u} - \mathbf{u}_{\mathcal{N}})(t, \cdot)\|^2 - \|(\mathbf{u} - \mathbf{u}_{\mathcal{N}})(t_m, \cdot)\|^2 + \nu \|\nabla(\mathbf{u} - \mathbf{u}_{\mathcal{N}})\|_{L^2(t_m, t; \mathcal{H})}^2 \leq \int_{t_m}^t \|R(\mathbf{u}_{\mathcal{N}})(t, \cdot)\|_{\nu^{-1/2}, -1}^2 dt.$$

Since $t \in (t_m, t_{m+1}]$ is arbitrary, we have proved the reliability (E.12). \square

Remark E.7 (Necessity of the “sufficient close” and the regularity conditions). *First, the main hurdle to prove the reliability is that the convection term $(*)$ is not positive definite, one would encounter some difficult in deriving the upper bound as yet moving $(*)$ to the right-hand side in (E.16) does not yield a meaningful estimate,*

$$\text{norm of the error} \simeq (I) + (II) = \langle R(\mathbf{u}_{\mathcal{N}}), \mathbf{u} - \mathbf{u}_{\mathcal{N}} \rangle - (*) \not\leq \langle R(\mathbf{u}_{\mathcal{N}}), \mathbf{u} - \mathbf{u}_{\mathcal{N}} \rangle$$

We note that Assumption (2.1.2) is equivalent to putting a threshold on $\|\mathbf{u}(t, \cdot)\|_{L^\infty}$, which is commonly assumed in the analysis of numerical methods for convection-dominated problems. A stronger alternative assumption than Assumption (2.1.2) would be imposing a stronger constraint on γ in Assumption (2.1.1) Using the “sufficient close” assumption (E.10) on the nonlinear term, we have

$$|(*)| \leq \|((\mathbf{u} - \mathbf{u}_{\mathcal{N}}) \cdot \nabla) \mathbf{u}_{\mathcal{N}}\|_{\mathcal{V}'} \|\mathbf{u} - \mathbf{u}_{\mathcal{N}}\|_{\mathcal{V}} \leq \gamma \|\nabla(\mathbf{u} - \mathbf{u}_{\mathcal{N}})\|^2.$$

Now if $\bar{\nu} := \nu - \gamma \geq 0$, one would reach a similar estimate as (E.17) by replacing ν with $\bar{\nu}$. If no extra assumption on γ is imposed.

Another way is to consider a Sobolev embedding directly after applying the Hölder inequality, for (E.15) we have

$$|(*)| \leq \int_{\Omega} |\nabla \mathbf{u}| |\mathbf{u} - \mathbf{u}_{\mathcal{N}}|^2 \leq \|\nabla \mathbf{u}\| \|\mathbf{u} - \mathbf{u}_{\mathcal{N}}\|_{L^4}^2 \leq \rho \|\nabla \mathbf{u}\| \|\nabla(\mathbf{u} - \mathbf{u}_{\mathcal{N}})\|^2,$$

where ρ is the constant in the following Sobolev embedding

$$\|\mathbf{u} - \mathbf{u}_{\mathcal{N}}\|_{L^4} \leq \rho \|\nabla(\mathbf{u} - \mathbf{u}_{\mathcal{N}})\|.$$

Now if $\tilde{\nu} := \nu - \rho \|\nabla \mathbf{u}\| \geq 0$, one would reach a similar estimate as (E.17) by replacing ν with $\tilde{\nu}$.

E.2 PROOF OF THEOREM 3.2

Theorem 2.2 (Equivalence of functional norm and negative Sobolev norm). *Consider the Gelfand triple $\mathcal{V} \Subset \mathcal{H} \Subset \mathcal{V}'$, then*

$$\|f\|_{\mathcal{V}'} = |f|_{-1} \text{ for } f \in \mathcal{H}/\mathbb{R}. \quad (\text{E.18})$$

Proof. This proof leverages the spectral basis for \mathcal{H} without the extra technicality of Schwartz space involved if one ought to assume that $f \in \mathcal{V}'$. Recall (3.1), and define an infinite version

$$\mathcal{S}_\infty := \text{span} \{e^{i\mathbf{k} \cdot \mathbf{x}} : \mathbf{k} \in 2\pi\mathbb{Z}^2\}. \quad (\text{E.19})$$

In what follows, we shall omit the extra technicality that one has to work on partial sums first, and then considers the convergence in the corresponding norms. We simply take for granted that the differentiation and integration/sum can be interchanged, and directly identify $f \in \mathcal{H}/\mathbb{R}$ with its Fourier series in \mathcal{S}_∞ ,

$$f(\mathbf{x}) = \sum_{\mathbf{k} \in 2\pi\mathbb{Z}^2} \hat{f}(\mathbf{k}) e^{i\mathbf{k} \cdot \mathbf{x}}.$$

Note since f lies in the quotient space, $\hat{f}(0, 0) = 0$. Then, the duality pairing can be identified as an L^2 -inner product, as well as the test function $v \in \mathcal{H}$,

$$\begin{aligned} \langle f, v \rangle_{\mathcal{V}', \mathcal{V}} &= (f, v) = \left(\sum_{\mathbf{k} \in 2\pi\mathbb{Z}^2} \hat{f}(\mathbf{k}) e^{i\mathbf{k} \cdot \mathbf{x}}, \sum_{\mathbf{m} \in 2\pi\mathbb{Z}^2} \overline{\hat{v}(\mathbf{m})} e^{i\mathbf{m} \cdot \mathbf{x}} \right) \\ &= \sum_{\mathbf{k} \in 2\pi\mathbb{Z}_n^2 \setminus \{0\}} \hat{f}(\mathbf{k}) \overline{\hat{v}(\mathbf{k})} \\ &\leq \left(\sum_{\mathbf{k} \in 2\pi\mathbb{Z}_n^2 \setminus \{0\}} |\mathbf{k}|^{-2} |\hat{f}(\mathbf{k})|^2 \right)^{1/2} \left(\sum_{\mathbf{k} \in 2\pi\mathbb{Z}_n^2 \setminus \{0\}} |\mathbf{k}|^2 |\hat{v}(\mathbf{k})|^2 \right)^{1/2} \\ &\leq |f|_{-1} |v|_1. \end{aligned}$$

As a result, by the definition in (3.7) and the estimate above

$$\|f\|_{\mathcal{V}'} = \sup_{v \in \mathcal{V}/\mathbb{R}, |v|_{\mathcal{V}}=1} |\langle f, v \rangle| = \sup_{v \in \mathcal{V}/\mathbb{R}} \frac{(f, v)}{|v|_1} \leq |f|_{-1}.$$

To show the other direction, let $v = (-\Delta)^{-1} f \in H^1(\mathbb{T}^2)$ by the well-posedness of $-\Delta v = f$ for $f \in L^2(\mathbb{T}^2)$. We have

$$\begin{aligned} \langle f, (-\Delta)^{-1} f \rangle &= (-\Delta v, v) = (f, (-\Delta)^{-1} f) \\ &= \left(\sum_{\mathbf{k} \in 2\pi\mathbb{Z}^2} \hat{f}(\mathbf{k}) e^{i\mathbf{k} \cdot \mathbf{x}}, \sum_{\mathbf{m} \in 2\pi\mathbb{Z}^2} (-\Delta_{\mathbf{x}})^{-1} \overline{\hat{f}(\mathbf{m})} e^{i\mathbf{m} \cdot \mathbf{x}} \right) \\ &= \left(\sum_{\mathbf{k} \in 2\pi\mathbb{Z}^2} \hat{f}(\mathbf{k}) e^{i\mathbf{k} \cdot \mathbf{x}}, \sum_{\mathbf{m} \in 2\pi\mathbb{Z}^2} \frac{1}{|\mathbf{m}|^2} \overline{\hat{f}(\mathbf{m})} e^{i\mathbf{m} \cdot \mathbf{x}} \right) \\ &= \sum_{\mathbf{k} \in 2\pi\mathbb{Z}_n^2 \setminus \{0\}} |\mathbf{k}|^{-2} |\hat{f}(\mathbf{k})|^2 = |f|_{-1}^2. \end{aligned}$$

Note that by the construction of v above, $|v|_1 = |f|_{-1}$, thus,

$$\sup_{v \in \mathcal{V}} \frac{(f, v)}{|v|_1} \geq \frac{(f, (-\Delta)^{-1} f)}{|v|_1} = |f|_{-1},$$

which proves the theorem. \square

E.3 CONVERGENCE RESULT OF FINE-TUNING

In the following theorem, Theorem E.9, we show that a sufficient condition for the optimizer to converge is to get in a neighborhood of the true solution, thus corroborating the necessity of training. Note in both Theorems 3.1 and E.9, the error term of $\mathbf{u}_{\mathcal{N}}^{(m)} - \mathbf{u}(t_m, \cdot)$ is present. We have to acknowledge that the “initial value” for the predicted trajectory, which is the last snapshot in the input trajectory, may have errors. This is the major motivation that we opt to use $\mathbf{u}_{\mathcal{N}}^{(\ell)}$, which is the input trajectory’s last snapshot in the skip-connection in \tilde{Q}

$$\mathbf{u}_{\mathcal{N}}^{(m+1)} = \mathbf{u}_{\mathcal{N}}^{(\ell)} + \tilde{Q}_{\theta}(\mathbf{v}_{\mathcal{N}}^{(m+1)}), \text{ for } m = \ell, \dots, \ell + n_t - 1. \quad (\text{E.20})$$

In some sense, SFNO learns the derivative $\partial_t \mathbf{u}$ ’s arbitrary-lengthed integral. If one wants to modify Algorithm 1 in view of Theorem E.9 such that the error control is guaranteed, the following algorithm can be used but loses the “parallel-in-time” nature. We also note that, thanks to the spectral convolution in \tilde{Q} being affine linear, showing Theorem E.9 is quite straightforward, as one has to establish the connection between fine-tuning and seeking a nonlinear Galerkin projection in Fourier space (3.1) under the functional norm. Let $\theta^* = \arg \min_{\theta} \|\mathcal{R}(\mathbf{u}_{\mathcal{N}}(\theta))(t_{m+1}, \cdot)\|_{-1}$, then $\tilde{Q}_{\theta^*}(\mathbf{u}_{\mathcal{N}}) = \arg \min_{\mathbf{v} \in \mathcal{S}} \|\mathbf{u}(t_{m+1}, \cdot) - (\mathbf{u}(t_m, \cdot) + \mathbf{v})\|_{*}$.

Algorithm 2 An error-guarantee fine-tuning strategy.

Input: ST-FNO $\mathcal{G}_{\theta, \Theta} := \tilde{Q}_{\theta} \circ \mathcal{G}_{\Theta}$; time stepping scheme $\mathcal{G}_{\alpha}(\cdot)$; optimizer $\mathcal{D}(\theta, \nabla_{\theta}(\cdot))$; training dataset: solution trajectories at $[t_1, \dots, t_{\ell'}]$ as input and at $[t_{\ell'+1}, \dots, t_{\ell'+n_t}]$ as output.

1: Train the ST-FNO model until the energy signature matches the inverse cascade.

2: **Freeze Θ in \mathcal{G}_{Θ} of ST-FNO $\mathcal{G}_{\theta, \Theta}$, keep \tilde{Q}_{θ} trainable.**

3: Cast all nn.Module involved and tensors to torch.float64 and torch.complex128 hereafter.

Input: Evaluation dataset: solution trajectories at $[t_1, \dots, t_{\ell}]$ as input, output time step n_t .

4: **for** $m = \ell, \dots, \ell + n_t - 1$ **do**

5: Extract the latent fields $\mathbf{v}_{\mathcal{N}}^{(m+1)}$ output of \mathcal{G}_{Θ} at t_{m+1} and hold them fixed.

6: By construction of ST-FNO: such that $\mathbf{u}_{\mathcal{N}}^{(m+1)}(\theta) := \mathbf{u}_{\mathcal{N}}^{(m)} + \tilde{Q}_{\theta}(\mathbf{v}_{\mathcal{N}}^{(m+1)})$.

7: March one step with $(\Delta t)^{\alpha}$ using \mathcal{G}_{α} : $D_t \mathbf{u}_{\mathcal{N}}(\theta) := (\Delta t)^{-\alpha}(\mathcal{G}_{\alpha}(\mathbf{u}_{\mathcal{N}}(\theta)) - \mathbf{u}_{\mathcal{N}}(\theta))$.

8: $j \leftarrow 0$

9: **while** $\eta_m(\mathbf{u}_{\mathcal{N}}(\theta), D_t \mathbf{u}_{\mathcal{N}}(\theta)) > \text{ToI}$ **do**

10: Apply the optimizer to update parameters in \tilde{Q} : $\theta \leftarrow \mathcal{D}(\theta, \nabla_{\theta}(\eta_m^2)), j \leftarrow j + 1$.

11: Forward pass only through \tilde{Q} to update $\mathbf{u}_{\mathcal{N}} \leftarrow \mathbf{u}_{\mathcal{N}}^{(m)} + \tilde{Q}_{\theta}(\mathbf{v}_{\mathcal{N}}^{(m+1)})$.

12: **if** $j > \text{Iter}_{\max}$ **then break**

13: $\mathbf{u}_{\mathcal{N}}^{(m+1)} \leftarrow \mathbf{u}_{\mathcal{N}}$

Output: A sequence of velocity profiles at corresponding time steps $\{\mathbf{u}_{\mathcal{N}}^{(m)}\}_{m=\ell+1}^{\ell+n_t}$.

Lemma E.8 (Local strict convexity for the fine-tuning loss). *Define $\|\cdot\|_{*,\delta}$ to be a (dual) graph norm on $L^{\infty}(\mathcal{T}_{\delta}; \mathbf{L}^2(\mathbb{T}^2)) \cap L^2(\mathcal{T}_{\delta}; \mathbf{H}^1(\mathbb{T}^2))$, where $\mathcal{T}_{\delta} := [t - \delta, t + \delta]$*

$$\|\mathbf{v}\|_{*,\delta} = \left\{ \int_{\mathcal{T}_{\delta}} \|\mathbf{v}(t, \cdot)\|^2 dt + \int_{\mathcal{T}_{\delta}} \|(\partial_t \mathbf{v} + \mathbf{v} \cdot \nabla \mathbf{v})(t, \cdot)\|_{\mathcal{V}'}^2 dt \right\}^{1/2}$$

For $\mathbf{u} \in L^{\infty}(\mathcal{T}_{\delta}; \mathbf{L}^2(\mathbb{T}^2)) \cap L^2(\mathcal{T}_{\delta}; \mathbf{H}^1(\mathbb{T}^2))$ the weak solution to (2.2) on \mathcal{T}_{δ} that is sufficiently smooth, there exists $\delta, \epsilon \in \mathbb{R}^+$, such that on

$$\mathcal{B}(\mathbf{u}; \epsilon) := \{\mathbf{v} \in L^{\infty}(\mathcal{T}_{\delta}; \mathbf{L}^2(\mathbb{T}^2)) \cap L^2(\mathcal{T}_{\delta}; \mathbf{H}^1(\mathbb{T}^2)) : \|\mathbf{u} - \mathbf{v}\|_{*,\delta} \leq \epsilon\},$$

the functional

$$J(\mathbf{v}(t, \cdot)) := \frac{1}{2} \|R(\mathbf{v})\|_{\mathcal{V}'}^2, \text{ where } R(\mathbf{v}) := \mathbf{f} - \partial_t \mathbf{v} - (\mathbf{v} \cdot \nabla) \mathbf{v} + \nu \Delta \mathbf{v}$$

is strictly convex.

Proof. First, by Theorem 3.2,

$$J(\mathbf{v}) = \frac{1}{2} \langle R(\mathbf{v}), (-\Delta)^{-1} R(\mathbf{v}) \rangle_{\mathcal{V}', \mathcal{V}}$$

Then,

$$\begin{aligned} DJ(\mathbf{v}; \boldsymbol{\xi}) &:= \lim_{\tau \rightarrow 0} \frac{J(\mathbf{v} + \tau \boldsymbol{\xi}) - J(\mathbf{v})}{\tau} = \frac{d}{d\tau} J(\mathbf{v} + \tau \boldsymbol{\xi}) \Big|_{\tau=0} \\ &= \langle DR(\mathbf{v}) \boldsymbol{\xi}, (-\Delta)^{-1} R(\mathbf{v}) \rangle_{\mathcal{V}', \mathcal{V}}. \end{aligned}$$

where

$$DR(\mathbf{v}) \boldsymbol{\xi} = \partial_t \boldsymbol{\xi} + \frac{1}{2} ((\boldsymbol{\xi} \cdot \nabla) \mathbf{v} + (\mathbf{v} \cdot \nabla) \boldsymbol{\xi}) - \nu \Delta \boldsymbol{\xi}$$

is the Fréchet derivative $DR(\mathbf{v}) : \mathcal{V} \rightarrow \mathcal{V}'$ by Lemma E.6. The Hessian is then

$$\text{Hess } J(\mathbf{v}; \boldsymbol{\xi}, \boldsymbol{\zeta}) = \langle DR(\mathbf{v}) \boldsymbol{\xi}, (-\Delta)^{-1} DR(\mathbf{v}) \boldsymbol{\zeta} \rangle + \langle \boldsymbol{\zeta} \cdot D^2 R(\mathbf{v}) \boldsymbol{\xi}, (-\Delta)^{-1} R(\mathbf{v}) \rangle.$$

Now, since $R(\mathbf{u}) = 0$ on \mathcal{T}_{δ} , we have

$$\text{Hess } J(\mathbf{u}; \boldsymbol{\xi}, \boldsymbol{\xi}) = \langle DR(\mathbf{v}) \boldsymbol{\xi}, (-\Delta)^{-1} DR(\mathbf{v}) \boldsymbol{\xi} \rangle.$$

If we assume that $DR(\mathbf{u})\xi \in \mathcal{V}'$ has its functional norm bounded above and below in $\mathcal{B}(\mathbf{u}; \epsilon)$, one has for any ξ

$$\|\xi\|_{\mathcal{V}}^2 \lesssim \|DR(\mathbf{u})\xi\|_{\mathcal{V}'}^2 \leq \text{Hess } J(\mathbf{u}; \xi, \xi).$$

Simply choosing ϵ small enough such that \mathbf{v} is sufficiently close to \mathbf{u} in graph norm associated with the PDE to make the coercivity above still true for $\text{Hess } J(\mathbf{v}; \xi, \xi)$ yields the desired local convexity. \square

Theorem E.9 (Guaranteed convergence of the fine-tuning). *In addition to the same assumptions with Theorem 3.1, suppose Lemma E.8 holds for $\epsilon \in (0, 1)$, and a given $\mathbf{u}_{\mathcal{N}}$ can be embedded in $\mathcal{B}(\mathbf{u}; \epsilon')$ for a $0 < \epsilon' \leq \epsilon$. Denote $\mathbf{u}_{\mathcal{N},j}^{(k)}$ the evaluation in Line 9 of Algorithm 2, and j the iteration of optimizer in Line 10, then the fine-tuning using the new loss function (3.9) produces a sequence $\{\mathbf{u}_{\mathcal{N},j}^{(k)}\}_{j=1}^{\infty} \subset \mathcal{B}(\mathbf{u}; \epsilon')$. Furthermore, suppose that the optimizer in Line 10 of Algorithm 2 has a learning rate converging to 0. Then, then the fine-tuning using the new loss function (3.9) produces a sequence of evaluations converging to the best possible approximation $\mathbf{u}_{\mathcal{N},\infty}^{(m+1)} \in \mathcal{S}$ of $\mathbf{u}(t_{m+1}, \cdot)$ starting from $\mathbf{u}_{\mathcal{N}}^{(m)}$, in the sense that for $m = \ell, \dots, \ell + n_t - 1$*

$$\|\mathbf{u}_{\mathcal{N},\infty}^{(m+1)} - \mathbf{u}(t_{m+1}, \cdot)\|_{\mathcal{V}} \leq \|\mathbf{u}_{\mathcal{N}}^{(l)} - \mathbf{u}(t_m, \cdot)\|_{\mathcal{V}} + c_1 n_t n^{-2} \|\mathbf{u}(t_{m+1}, \cdot)\|_2 + c_2 n_t (\Delta t)^{\alpha-1}. \quad (\text{E.21})$$

Proof. For simplicity, we denote $\mathbf{u}_0 := \mathbf{u}(t_l, \cdot) \in \mathcal{V}$, and $\mathbf{u}_{\mathcal{N}} := \mathbf{u}_{\mathcal{N}}^{(m)} \in \mathcal{S}|_{t=t_m} =: \mathcal{S}$. First, by line 10 fine-tuning algorithm, if one solves the optimization in the functional norm exactly, we have

$$\theta^* = \arg \min_{\theta} \delta t \|\bar{R}(\mathcal{G}_{\alpha} \mathbf{u}_{\mathcal{N}}(\theta))\|_{\mathcal{S}'}^2, \quad \text{where } \mathbf{u}_{\mathcal{N}}(\theta) := \mathbf{u}_{\mathcal{N}} + \tilde{Q}_{\theta}(\mathbf{v}_{\mathcal{N}}) \quad (\text{E.22})$$

where $\bar{R}(\cdot)$ is the residual functional computed using the time derivative term from an extra-fine-step solver's result

$$\bar{R}(\mathbf{v}) := \mathbf{f} - D_t \mathbf{v} - (\mathbf{v} \cdot \nabla) \mathbf{v} + \nu \Delta \mathbf{v}$$

Unlike representing the derivative using output from the neural operator, one of the keys of our algorithm is that the error for $\partial_t \mathbf{u} - D_t \mathbf{u}_{\mathcal{N}}(\theta)$ can be explicitly estimated using the framework to develop estimates for truncation error in traditional time marching schemes for NSE. Due to the choice of the time step, this truncation error will be of higher order. Note for any $\mathbf{v}_{\mathcal{N}}, \tilde{Q}_{\theta}(\mathbf{v}_{\mathcal{N}}) \in \mathcal{S}$, as a result, solving (E.22) is equivalent to solve the following: denote $\mathbf{u}_{\delta} := \mathcal{G}_{\alpha}(\mathbf{u}_{\mathcal{N}} + \delta \mathbf{u})$ for $\delta \mathbf{u} \in \mathcal{S}$

$$\min_{\delta \mathbf{u} \in \mathcal{S}} (\delta t)^{1/2} \|D_t(\mathbf{u}_{\delta}) + (\mathbf{u}_{\delta} \cdot \nabla) \mathbf{u}_{\delta} - \nu \Delta \mathbf{u}_{\delta} - \mathbf{f}\|_{\mathcal{S}'}.$$

Replacing $D_t(\cdot)$ by $\partial_t(\cdot)$ we have a truncation error term, whose error is of order $(\Delta t)^{\alpha-1}$ due to taking the time derivative:

$$\begin{aligned} & \|D_t(\mathbf{u}_{\delta}) + (\mathbf{u}_{\delta} \cdot \nabla) \mathbf{u}_{\delta} - \nu \Delta \mathbf{u}_{\delta} - \mathbf{f}\|_{\mathcal{S}'} \\ & \leq \|\partial_t \mathbf{u}_{\delta} + (\mathbf{u}_{\delta} \cdot \nabla) \mathbf{u}_{\delta} - \nu \Delta \mathbf{u}_{\delta} - \mathbf{f}\|_{\mathcal{S}'} + \|D_t(\mathbf{u}_{\delta}) - \partial_t \mathbf{u}_{\delta}\|_{\mathcal{S}'} \end{aligned}$$

We focus on estimating the first term above,

$$\min_{\delta \mathbf{u} \in \mathcal{S}} (\delta t)^{1/2} \|\partial_t \mathbf{u}_{\delta} + (\mathbf{u}_{\delta} \cdot \nabla) \mathbf{u}_{\delta} - \nu \Delta \mathbf{u}_{\delta} - \mathbf{f}\|_{\mathcal{S}'} \quad (\text{E.23})$$

By the fact that $\|\cdot\|_{\mathcal{S}'}$ inherit the scaling law and the triangle inequality from $\|\cdot\|_{\mathcal{V}'}$, it is convex as well in this neighborhood. As a result, any gradient-based optimizer with a converging step size shall converge to the minimum, achieved at $\mathbf{u}_{\mathcal{N},\infty}^{(m)}$, with a linear convergence rate. Moreover, $\mathbf{u}_{\mathcal{N},\infty}^{(m)} := \mathbf{u}_{\mathcal{N}}^{(m)} + (\delta \mathbf{u})^*$ is the (nonlinear) Galerkin projection of $\mathbf{u}(t_m, \cdot) \in \mathcal{V}$. \square