Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

SIMR-NO: A Spectrally-Informed Multi-Resolution Neural Operator for Turbulent Flow Super-Resolution

Muhammad Abid

mabid@vols.utk.edu

University of Tennessee at Knoxville

Omer San

University of Tennessee at Knoxville

Research Article

Keywords:

Posted Date: November 11th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11514570/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

SIMR-NO: A Spectrally-Informed Multi-Resolution Neural Operator for Turbulent Flow Super-Resolution

Muhammad Abid[1*] and Omer San[1]

[1]Department of Mechanical and Aerospace Engineering, University of Tennessee, Knoxville, 37996, TN, USA.

*Corresponding author(s). E-mail(s): mabid@vols.utk.edu;
Contributing authors: osan@utk.edu;

Abstract

Reconstructing high-resolution turbulent flow fields from severely under-resolved observations is an ill-posed inverse problem of central importance in computational fluid dynamics. Classical interpolation cannot recover the missing fine scales, and convolutional super-resolution networks lack spectral and multiscale inductive biases suited to turbulence at large upscaling factors. We introduce the *Spectrally-Informed Multi-Resolution Neural Operator* (SIMR-NO), a hierarchical operator-learning framework that factorizes the reconstruction across intermediate resolutions. Each stage combines a deterministic interpolation prior with a residual correction computed by Fourier layers whose spectral multipliers are modulated by a learned radial gate, and local convolutional refinement modules recover features beyond the truncated Fourier basis. SIMR-NO is evaluated on two two-dimensional flows with distinct dynamics, statistically stationary Kolmogorov-forced turbulence and freely decaying homogeneous isotropic turbulence at Reynolds numbers of **1000**, **4000**, and **8000**, in which $\mathbf{128 \times 128}$ vorticity fields are reconstructed from $\mathbf{8 \times 8}$ observations, a $\mathbf{16\times}$ upscaling factor. In comparison with bicubic interpolation, the Fourier neural operator, EDSR, and LapSRN, SIMR-NO achieves the lowest reconstruction error and the highest structural similarity for both flows and at every Reynolds number. On decaying turbulence, SIMR-NO attains a mean relative error of about **3%** at every Reynolds number, up to **55%** lower than that of the strongest baseline, and its accuracy is nearly insensitive to the Reynolds number. Spectral diagnostics, vorticity-conditioned error statistics, and localized comparisons show that these gains are accompanied by a more accurate reconstruction of intense vortex cores, fine-scale filaments, and high-wavenumber content.

Keywords: Scientific machine learning, Neural operators, Turbulence super-resolution, Multiscale modeling, Spectral learning

1 Introduction

The process of reconstructing high-resolution turbulent flow fields from extremely limited observational data serves as an essential inverse problem that connects the fields of computational fluid dynamics (CFD) and scientific machine learning (SciML) research. The practical world only permits researchers to observe flow patterns through spatial measurements, which are too broad to be useful for their numerical simulations and laboratory experiments and field measurements because their equipment does not provide sufficient resolution [1]. The field of fluid dynamics uses various physical quantities, which are required to study fine-scale structures that include vortical filaments and shear layers and small-scale turbulent fluctuations. These structures control all three turbulent dynamic processes, which include energy transfer and mixing and dissipation processes that are central to understanding turbulent dynamics [2]. The process of recovering missing spatial scales from restricted observations happens to be more than just a technical feature because it serves as an essential requirement that need to perform accurate analyses and predictions about turbulent systems and their corresponding models [3–5].

The solution to this problem becomes challenging because downsampling leads to permanent data loss. The high-resolution flow field becomes inaccessible because only coarse-resolution data provides visibility to high-frequency spatial details that create an inverse reconstruction problem that is ill-posed in the sense of Hadamard: many high-resolution fields are consistent with the same coarse data, and the reconstruction is sensitive to perturbations of the input [6]. Nonlinear multiscale interactions in turbulent flows aggravate this difficulty, because the cascade redistributes energy and enstrophy across scales, so that the unresolved scales are dynamically coupled to, but not determined by, the resolved ones [7, 8]. Classical interpolation techniques such as bilinear or bicubic interpolation show the ability to reconstruct smooth large-scale variations, but they cannot recover the small-scale vortical structures that create turbulence and must follow physical laws. The resulting reconstructions become excessively smooth because they constantly misrepresent fine-scale energy components that control important downstream quantities.

The scientific field of machine learning has developed new techniques that use data to create effective methods for studying intricate physical systems. The method of operator learning provides a structured approach that enables researchers to create mathematical models that describe how infinite-dimensional mappings between function spaces behave under partial differential equations (PDEs). The operation of neural operators differs from traditional neural networks because they learn to create continuous mappings between function spaces, which allow them to make predictions without needing specific dimensional inputs [9]. DeepONet [10] and the Fourier Neural Operator (FNO) [11] have established themselves as essential frameworks that enable the development of mathematical models that simulate various physical phenomena,

including fluid dynamics, climate change, and elasticity [12]. A main benefit of these models is that their parameterization is defined independently of the discretization, so that a trained operator can be evaluated at different resolutions, and that spectral layers represent nonlocal interactions efficiently. Researchers have developed neural operator systems that now support physical constraints and irregular geometries while using advanced transformer and multiscale systems [13–17] which now enable the study of complicated scientific problems.

Simultaneously, machine learning methods have increasingly been applied to turbulence modeling and flow reconstruction. Data-driven approaches have been developed for turbulence closure modeling [18], subgrid-scale modeling [19], reduced-order modeling [20], and machine-learning-accelerated CFD [3–5, 21], collectively demonstrating the transformative potential of learning-based methods in computational fluid mechanics. More specifically, a growing body of work has demonstrated that deep neural networks can reconstruct high-resolution turbulent flow fields from coarse observations with substantially improved accuracy relative to classical interpolation baselines [22–27]. The methods depend on convolutional neural network (CNN) architectures, which were developed to solve image super-resolution problems while using residual learning and dense connectivity and adversarial training to enhance visual quality. The models successfully capture local spatial relationships; however, they fail to recognize the fundamental spectral and multiscale patterns that define turbulent flows. The limitation becomes more significant when using higher upscaling factors because they need to restore accurate fine-scale details, which presents the most difficult task.

Despite significant advances in super-resolution, reconstructing high-resolution turbulent flow fields from coarse observations remains challenging, particularly when large-scale structures and fine-scale fluctuations must be recovered simultaneously. Conventional single-stage approaches typically learn a direct mapping from coarse inputs to high-resolution outputs, which can become increasingly difficult as the resolution gap grows [28, 29]. Previous studies have also explored multistage reconstruction strategies, including the two-stage framework of [30], demonstrating that progressive refinement is not unique to the present work. Nevertheless, accurately recovering multiscale turbulent structures and their associated physical characteristics remains an important challenge.

For turbulent flows, reconstruction quality cannot be assessed solely through pointwise field accuracy, since physically relevant information is also reflected in energy spectra, enstrophy distributions, modal energy content, and coherent vortical structures [31–33]. Consequently, a super-resolution model may achieve a low mean squared error while exhibiting discrepancies in spectral energy distributions or vortical topology, potentially limiting its reliability for subsequent flow analysis and modeling [34]. These considerations motivate the development of spectrally informed hierarchical operator-learning frameworks that combine progressive multiresolution refinement with the reconstruction of physically meaningful turbulent flow structures.

The current study develops a new architectural design, which it names *Spectrally-Informed Multi-Resolution Neural Operator* (SIMR-NO), to solve existing challenges. The core idea is to reformulate turbulence super-resolution as a hierarchical inverse operator learning problem, decomposing the reconstruction operator into a sequence

of progressively refined stages rather than recovering all missing spatial scales through a single monolithic mapping. At each stage, a deterministic bicubic interpolation prior provides a smooth structural estimate, upon which the learned neural operator applies a spectral residual correction to restore lost high-frequency content. Each stage doubles the spatial resolution of its input, so that the reconstruction proceeds incrementally from coarse to fine scales [35, 36]. This stands in direct contrast to single-stage architectures, which must simultaneously recover all spatial scales through one highly nonlinear mapping, and to U-Net-type architectures, which couple encoder-decoder representations at fixed resolution levels and lack flexibility to generalize across arbitrary scales. The hierarchical formulation instead reduces the upscaling factor that each learned stage must bridge from 16× to 2×, connecting to the well-established advantages of classical multigrid solvers [37] and progressive training strategies [38, 39], in which coarse-to-fine decompositions are used to accelerate convergence.

The primary architectural advancement of SIMR-NO structural design emerges through its implementation of spectrally gated convolution layers, which manage learned Fourier multipliers according to their radial wavenumber. In addition to the mode-wise multipliers of standard FNO layers, a learned function of the radial wavenumber modulates all retained modes, providing an explicit, shared weighting of spectral corrections by scale. The energy distribution of turbulent flows follows low-frequency patterns while their enstrophy moves to higher-frequency patterns according to [31, 32, 40]. The combination of hierarchical reconstruction with spectrally informed operator layers establishes a unified framework in SIMR-NO, which connects operator learning, multiscale decomposition, and spectral modeling methods to achieve accurate turbulence super-resolution results.

The proposed architecture is evaluated on two complementary benchmarks of two-dimensional turbulence with distinct dynamics. The first is Kolmogorov-forced turbulence, a statistically stationary flow sustained by a steady large-scale forcing [27, 32, 41]. The second is decaying homogeneous isotropic turbulence (DHIT), an unforced flow in which energy and enstrophy decay freely and the spectral content of the field evolves continuously in time [41, 42], which we consider at three Reynolds numbers, $Re = 1000$, 4000, and 8000. In both cases, 128×128 vorticity fields are reconstructed from extremely coarse 8×8 observations, which represents a 16× upscaling factor. At this ratio the coarse grid resolves only wavenumbers $k \leq 4$, so that most of the vortical structure must be inferred rather than interpolated. The main contributions of this work are summarized as follows:

- We formulate turbulence super-resolution as a multi-stage inverse problem that progressively reconstructs the field across intermediate resolutions.
- We introduce SIMR-NO, a multi-resolution neural operator that combines interpolation priors with spectrally informed residual corrections at each stage.
- We propose a learned radial-wavenumber gate that weights the spectral corrections of Fourier layers by scale.
- We evaluate SIMR-NO on stationary Kolmogorov-forced turbulence and on decaying homogeneous isotropic turbulence at $Re = 1000$, 4000, and 8000.

- SIMR-NO outperforms bicubic interpolation, FNO, EDSR, and LapSRN in all cases. It recovers high-wavenumber content and intense vortex cores more accurately, and its advantage grows with the Reynolds number.

The remainder of the paper is organized as follows. Section 2 reviews related work on neural operators and turbulence super-resolution. Section 3 describes the physical dataset and the mathematical formulation of the reconstruction problem. Section 4 presents the SIMR-NO architecture in detail. Section 5 describes the experimental setup and evaluation protocol. Section 6 reports quantitative and qualitative reconstruction results. Finally, Section 7 concludes the paper and outlines directions for future work.

2 Related Work

2.1 Neural Operators and Operator Learning

Operator learning serves as an effective method for creating mathematical models that describe infinite-dimensional mappings between function spaces. The primary goal of this research is to develop a unified operator that can handle all different variations of a partial differential equation. The approach used here differs from standard neural networks, which require separate training for every new input because they process fixed-dimensional vector data. By learning mappings between function spaces directly, neural operators enable discretization-agnostic inference and resolution-independent generalization [9], properties that are especially valuable in scientific computing where simulation data may be available at multiple resolutions and where generalization to unseen discretizations is essential for practical deployment [12].

The first functional method that developed this field of research was DeepONet [10], which created the necessary theoretical basis to use its branch-trunk network structure for neural operator approximation. The method showed that operators can be accurately approximated by using computational methods that follow the universal approximation theorem, thus enabling operators to be learned from data in various types of PDE systems. The Fourier Neural Operator (FNO) developed by [11] establishes a scalable system that uses Fourier domain parameters to create global convolution kernel functions. FNO models use Fourier multipliers to train their systems, which enables them to model long-range spatial dependencies and nonlocal interactions through truncated spectral modes. The model shows its highest performance level when it handles physical systems that use elliptic or parabolic PDEs, which produce smooth solutions. The FNO model, together with its different variations, has demonstrated successful performance across several forward modeling tasks, which include simulating Navier-Stokes flows and estimating Darcy permeability and predicting global weather patterns [43].

Researchers established a deeper understanding of neural operators through their studies, which used existing theoretical foundations as their basis for development. Kovachki et al. [12] developed a mathematical framework that integrates all operator learning approximation techniques used in Banach spaces, while the framework specifies the precise conditions that allow neural operators to model all continuous

nonlinear operators. Theoretical progress has established how the adjoint function operates within operator learning [44]. The research shows precise algorithmic performance metrics, which include both convergence rates and generalization limits [45]. The recent advancements in neural operators have established their mathematical foundation, which matches the standards of traditional numerical techniques used in finite element analysis and spectral methods. The established foundation identifies its correct application for essential scientific research.

The developed architectural extensions for neural operators, which now provide enhanced flexibility together with improved capabilities to manage complex physical systems. The deformation-based FNO [14] operator learning system performs irregular domain learning through its ability to learn deformation patterns that transform all geometric shapes into standard grid formats. The study of physics-informed neural operators [13] uses partial differential equation residuals as soft training constraints, which help them achieve better physical accuracy while learning from small datasets. U-shaped neural operator architectures [15] use encoder-decoder structures with skip connections to create a system that automatically extracts multiscale features from different spatial levels. The GNOT operator architecture from [16] uses attention mechanisms to model nonlocal dependencies that exist between different physical domains and unstructured query points. Improved FNO variants [17] have addressed spectral aliasing and expressiveness limitations of the original architecture through modified normalization and operator parameterizations. The use of neural operators for inverse reconstruction tasks, which require recovering hidden detailed fields from incomplete coarse observations, remains less developed than its application in forward modeling. The current research specifically aims to fill this research gap through its examination of turbulence super-resolution, which serves as a fundamental and essential practical example of inverse problems.

2.2 Machine Learning for Turbulence Reconstruction

Researchers have started using machine learning techniques for turbulence research, which includes turbulence closure models and reduced-order models and flow state estimation and data-driven reconstruction methods. The work of Ling et al. [18] showed that neural networks that contain Galilean invariant features can predict Reynolds stress anisotropy tensors with greater accuracy than standard linear eddy-viscosity closure models. The study by Maulik and San [19] demonstrated that two-dimensional turbulence modeling could be improved through the use of convolutional networks, which developed subgrid-scale models that surpassed traditional Smagorinsky-type eddy-viscosity methods. The research work after this conducted an in-depth assessment of architectural design alternatives, which studied their impact on physical accuracy preservation during different turbulent flow situations [46]. Beck et al. [47] created a three-dimensional large-eddy simulation through data-driven closure methods, which they used with deep residual networks to demonstrate that learned subgrid models could function as substitutes for conventional models. The training process of learned turbulence models now achieves better physical consistency through physics-informed methods, which use governing equation residuals as training objectives [48]. The researchers demonstrated that data-driven methods could accurately model

turbulence through their collaborative work, which resulted in the advancement of machine learning-based CFD acceleration research [3–5, 21].

The research field has grown increasingly active because researchers use deep learning techniques to perform turbulence super-resolution, which reconstructs high-resolution flow fields from coarse observations. The research by Fukami et al. [22] proved that convolutional neural networks trained on paired coarse-to-fine flow data can achieve better results than bicubic interpolation when recovering high-resolution turbulent flow structures from low-resolution inputs. Their work established a benchmark pipeline and evaluation methodology that has since been widely adopted and extended by the community. The subsequent research demonstrated that spatial-temporal super-resolution evaluation required deep learning models to retrieve complete turbulent dynamic information, which was restricted to spatial and temporal measurements [23]. The single-snapshot setting [49, 50] addresses the practically important scenario where time-series data are unavailable, which requires the model to create detailed structural reconstructions from only one basic observation that lacks time information. The comprehensive review presents contemporary research developments in machine learning-based flow reconstruction and discovers existing research gaps in the field[25, 26, 51].

The advancement of turbulence super-resolution methods has advanced through two main approaches which include physics-based methods and generative techniques. Gao et al. [52] developed physics-informed CNNs which employ Navier-Stokes residuals to establish training loss functions that enable super-resolution without requiring matched high-resolution training data. The study from [53] shows that unsupervised methods can learn super-resolution mappings by using statistical characteristics of turbulent flow patterns without needing high-resolution paired data for training. Researchers have developed graph neural network architectures that [54–56] enable flow reconstruction through irregular and unstructured mesh systems, which provide more geometric flexibility than structured convolutional methods. The research area now investigates three-dimensional turbulence reconstruction, which uses two-dimensional observation data [57]while studying subgrid-scale super-resolution for large-eddy simulation pipelines [58–62] and implementing loss functions that address spectral and dynamical discrepancies between reconstructed and target fields [63, 64]. The probabilistic and diffusion-based methods work for uncertainty estimation in super-resolved fields by delivering confidence assessments that accompany their point reconstruction outputs [65]. This feature serves important benefits that assist in analyzing turbulent flow applications.

Most current turbulence super-resolution solutions exhibit a fundamental design constraint, which requires them to use convolutional neural networks that were developed for image restoration work. These networks use unchanged spatial filtering windows to detect nearby spatial patterns in data, but they fail to represent the complete atmospheric disturbance behavior that exists in turbulent flow. Standard CNNs have no built-in mechanism to distinguish energy-containing large-scale modes from enstrophy-cascading small-scale modes [31, 32], nor to adaptively emphasize spectral corrections at physically relevant wavenumber ranges. The limitation becomes more important when the system applies large upscaling factors because the system must

reconstruct accurate fine-scale structures instead of creating smooth large-scale variations. The standard mean squared error training objectives fail to penalize spectral mismatches, which enables a model to achieve low pixel-level error while producing physically incorrect energy distributions [27, 34]. The present work directly addresses both of these limitations by embedding spectral awareness into the reconstruction operator through Fourier-domain gating mechanisms and by adopting a hierarchical reconstruction strategy that progressively recovers fine-scale content across spatial scales.

2.3 Deep Learning for Image Super-Resolution

The research field of image super-resolution functions as the fundamental basis that develops scientific super-resolution methods through its provision of architectural designs and training methods and evaluation standards that are used in the field of fluid mechanics. The SRCNN [28] system proved that end-to-end learned super-resolution systems based their operation on shallow convolutional networks. The system achieved significant performance gains because its three-layer design processed paired low- and high-resolution image patches during training. The researchers conducted their studies to build deeper networks, which achieved better results through the application of residual learning methods from [66] and the combination of increased network depth with skip connections from [36]. The researchers used recursive unfolding techniques from [67, 68] to enhance parameter efficiency because the method allowed weight sharing across different reconstruction processes. The research used two different methods of sub-pixel upsampling, which improved its inference performance through efficient feature extraction in low-resolution space followed by one learned upsampling process [69, 70].

The present research work utilizes multiple architectural designs that were developed during this historical time frame because these designs use multiscale and progressive design methods. The system created by LapSRN [35] uses a Laplacian pyramid for progressive upsampling, which allows the system to reconstruct an image through multiple resolution stages instead of processing it at once. The pyramid system contains multiple levels, each of which uses a small convolutional branch to calculate the difference between the current upscaled prediction and the actual result that exists at that particular level. The organization of this system into hierarchical structures enables users to learn large-magnification super-resolution because the system breaks down complex tasks into easier components that directly support the multiple-resolution decomposition method used in SIMR-NO. EDSR [29] proved that deep residual networks achieve better performance without batch normalization because the normalization layers remove essential scale information needed for accurate reconstruction. EDSR serves as a strong convolutional baseline in our experiments. RDN [71] and DBPN [72] used dense connectivity and iterative back-projection mechanisms to enhance feature reuse and multi-scale representation learning, which resulted in better recovery of high-magnification fine-grained structures.

The development of image super-resolution has progressed through the recent introduction of attention-based and generative approaches. RCAN [73] channel attention

networks, proved that using feature channel weights according to their information value enables the creation of high-quality reconstruction results that efficiently recover detailed textures. The ESRGAN [74] adversarial approach changed its optimization process to focus on training a discriminator that distinguishes between super-resolved and actual high-resolution images, which resulted in sharper visual output but decreased pixel accuracy. The SwinIR [75] transformer restoration model used self-attention mechanisms from [76] to develop advanced reconstruction capabilities that use long-distance spatial information that convolutional systems with restricted field of view limitations cannot track. The diffusion model-based super-resolution method in [77] introduces stochastic iterative refinement as a systematic approach that produces multiple accurate high-quality reconstructions that show the uncertainty of high-frequency details, which is essential for turbulence studies that deal with unpredictable fine-scale patterns that can only be observed through basic measurements. The efficient cascaded architectures from [78] together with the task-adaptive frameworks from [79], enable achieving better super-resolution results under their current system resource constraints while solving actual implementation challenges.

Architectural advancements introduce essential design concepts, which include hierarchical reconstruction and residual correction and attention-based feature weighting and long-range dependency modeling. The process of turbulence super-resolution operates through a fundamental distinction from image restoration because it requires physical correctness to achieve proper results. The reconstructed vorticity field displays a clear appearance yet fails to meet perceptual accuracy because it fails to generate the correct energy spectrum and enstrophy distribution together with the specific vortical structure [34]. Standard image quality metrics such as SSIM and PSNR measure structural and luminance similarity, but they do not assess spectral inconsistencies or modal energy errors or violations of the governing conservation principles. The scientific super-resolution community has started to acknowledge the fundamental misalignment between image restoration objectives and physical fidelity requirements because this issue directly affects their research work [30, 63, 64]. The present study develops architectural designs that use spectral inductive biases to model turbulent flow physics according to its natural behavior.

2.4 Positioning of the Present Work

The SIMR-NO framework combines three different research areas, which include neural operator learning and turbulence super-resolution and hierarchical image reconstruction. The system uses common research methods to build a system that unifies their different strengths while solving their separate research problems. The SIMR-NO system uses a hierarchical multi-resolution framework that enables the system to resolve the inverse mapping problem through multiple spatial resolution levels. FNO systems use standard architecture designs to handle forward modeling tasks that require matching input data with output data that uses similar resolution grids. The architecture of FNO systems uses one operational design to handle tasks that require gradual recovery of detailed elements that appear in high-definition super-resolution. The hierarchical decomposition in SIMR-NO reduces the effective learning complexity at each

stage, which results in better optimization conditions because the system divides difficult global mapping into simpler residual corrections. The system enables gradual development of multiscale corrections that match the turbulent energy cascade [7, 8].

The architectural design of SIMR-NO, which uses neural operator systems and their Fourier-domain gating system for spectral data processing, exists as a distinct approach to super-resolution modeling when compared with EDSR and LapSRN. The LapSRN method follows a method of progressive reconstruction, yet its entire system depends on local convolutional processes that use unchanging receptive field dimensions, which fail to capture the complete spectral characteristics of turbulent flow patterns. SIMR-NO replacement of convolutional residual blocks at all pyramid levels uses spectrally gated Fourier operator layers that provide nonlocal scale-aware corrections that match the physical wavenumber-dependent statistical behavior of two-dimensional turbulence [32, 41]. SIMR-NO uses built-in physical architectural features to create its system design, but physics-informed methods use soft penalty terms to implement their PDE constraints in training loss functions according to [13, 48, 52]. The spectrally gated convolution system provides precise frequency-based correction capabilities that match the two-dimensional turbulence spectral characteristics of its dual cascade process that separates energy at large scales and enstrophy at small scales [31, 40] because it does not need access to the governing Navier-Stokes equations during training. The SIMR-NO system functions in situations where practitioners lack knowledge of the base PDE and face challenges to solve it and when experimental data replaces numeric simulation data.

SIMR-NO functions as a deterministic operator learning system that operates according to established scientific principles while maintaining accurate spectral measurements, and it does not emulate random behavior [65, 77]. While generative models can create multiple possible outcomes that match their basic data input, they need higher processing power during their testing phase because they do not maintain spectral consistency with actual turbulence patterns. SIMR-NO provides a fast single-pass reconstruction process that effectively maintains both energy spectra and enstrophy distribution, thus making it suitable for use in large data assimilation and simulation speed enhancement tasks. The design choices of SIMR-NO create a system that combines physical principles with architectural design to achieve better turbulence super-resolution results through its improved reconstruction accuracy and better physical performance than current methods across all tested conditions. The research results in Sections 4 and 6 demonstrate that this location between two positions results in ongoing improvements which affect standard reconstruction metrics and spectral and modal diagnostic methods used to evaluate physical performance.

3 Dataset and Problem Formulation

This section describes the physical systems, the datasets, and the mathematical formulation of the turbulence super-resolution problem studied in this work. Two configurations of two-dimensional incompressible turbulence are considered. The first is Kolmogorov-forced turbulence, a statistically stationary flow maintained by a steady

large-scale body force, which is a standard benchmark for data-driven flow reconstruction because of its broad range of active scales and well-established statistical properties [31, 32]. The second is decaying homogeneous isotropic turbulence (DHIT), an unforced flow evolving from a random initial condition, which is a canonical test case for numerical methods and subgrid models in two-dimensional turbulence [41, 42]. Both flows obey the same governing equations and are reconstructed with the same measurement model and operator-learning formulation; they differ in forcing, stationarity, and the distribution of vortical energy across scales.

3.1 Governing Equations and Physical Setup

The underlying physical system is the two-dimensional incompressible Navier–Stokes equations written in vorticity–streamfunction form. Let $\Omega = (0, 2\pi)^2$ denote a periodic spatial domain, and let $\omega : \Omega \times [0, T] \to \mathbb{R}$ denote the scalar vorticity field. The evolution of ω is governed by

$$\frac{\partial \omega}{\partial t} + \mathbf{u} \cdot \nabla \omega = \nu \Delta \omega + f(x, y), \qquad (x, y) \in \Omega, \quad t \in [0, T], \tag{1}$$

where $\nu > 0$ denotes the kinematic viscosity and $f : \Omega \to \mathbb{R}$ is the curl of an external body force. The incompressible velocity field $\mathbf{u} = (u, v) : \Omega \times [0, T] \to \mathbb{R}^2$ satisfies the divergence-free constraint $\nabla \cdot \mathbf{u} = 0$ and is recovered from the vorticity through the streamfunction $\psi : \Omega \times [0, T] \to \mathbb{R}$ via the kinematic relations

$$\mathbf{u} = \nabla^{\perp} \psi \coloneqq (-\partial_y \psi,\ \partial_x \psi), \qquad \omega = \Delta \psi, \tag{2}$$

where the second relation is the vorticity-streamfunction Poisson equation. Given ω, the streamfunction is recovered by solving $\Delta \psi = \omega$ subject to periodic boundary conditions, which is accomplished efficiently in Fourier space via

$$\widehat{\psi}_{k_x, k_y} = -\frac{\widehat{\omega}_{k_x, k_y}}{k_x^2 + k_y^2}, \qquad (k_x, k_y) \neq (0, 0), \tag{3}$$

and $\widehat{\psi}_{0,0} = 0$ by convention; the zero mode of ω vanishes because ω is the curl of a periodic velocity field. The velocity field is then obtained by applying $\nabla^{\perp}$ in Fourier space, which satisfies the incompressibility constraint exactly for the resolved modes.

3.1.1 Kolmogorov-forced turbulence

To sustain a statistically stationary turbulent state, the flow is driven by the Kolmogorov body force $\chi \sin(k_f y)\, \mathbf{e}_x$ with amplitude $\chi = 1$ and wavenumber $k_f = 4$ [32, 80]. Its curl gives the vorticity forcing of Eq. (1),

$$f(x, y) = -\chi\, k_f \cos(k_f y), \qquad k_f = 4, \tag{4}$$

which injects energy at the single wavenumber k_f and produces the characteristic large-scale structures of the vorticity field; in the statistically stationary state, the

energy input is balanced on average by viscous dissipation. Following [80], the Reynolds number on the domain $(0, 2\pi)^2$ is

$$Re = \frac{\sqrt{\chi}}{\nu} = \frac{1}{\nu} = 40, \tag{5}$$

that is, $\nu = 1/40$.

3.1.2 Decaying homogeneous isotropic turbulence

In the second configuration no energy is injected. The flow is initialised with a broadband random vorticity field and evolves freely under nonlinear transfer and viscous dissipation; it is therefore statistically homogeneous and isotropic but non-stationary. The governing equations are Eqs. (1)–(3) with $f \equiv 0$ on the periodic domain $\Omega = (0, 2\pi)^2$, which in nondimensional vorticity–streamfunction form read

$$\frac{\partial \omega}{\partial t} + J(\psi, \omega) = \frac{1}{Re}\,\Delta\omega, \qquad \Delta\psi = \omega, \qquad (x, y) \in \Omega,\ t \in [0, T], \tag{6}$$

where $J(\psi, \omega) = \mathbf{u} \cdot \nabla\omega$ is the advection term and $Re = 1/\nu$. Lengths are scaled such that the domain has size 2π and velocities such that the initial kinetic energy equals $1/2$.

With the kinetic energy, enstrophy, and palinstrophy

$$K(t) = \tfrac{1}{2}\langle|\mathbf{u}|^2\rangle, \qquad Z(t) = \tfrac{1}{2}\langle\omega^2\rangle, \qquad P(t) = \tfrac{1}{2}\langle|\nabla\omega|^2\rangle, \tag{7}$$

where $\langle\cdot\rangle$ denotes the average over Ω, Eq. (6) yields the exact budgets

$$\frac{dK}{dt} = -\frac{2}{Re}\,Z, \qquad \frac{dZ}{dt} = -\frac{2}{Re}\,P, \tag{8}$$

so both quantities decrease in time. Energy is transferred toward large scales and enstrophy toward small scales, where it is dissipated [31, 32, 42]. As a result, the spectral peak moves to larger scales and each snapshot represents a distinct flow state.

Following [41], the initial energy spectrum is

$$E(k, 0) = \frac{4}{3\sqrt{\pi}}\,\frac{k^4}{k_0^5}\exp\left[-\left(\frac{k}{k_0}\right)^2\right], \qquad k = \|\mathbf{k}\|_2, \tag{9}$$

with peak wavenumber $k_0 = 10$ and $\int_0^\infty E(k, 0)\,dk = 1/2$. The initial vorticity assigns this spectrum to each Fourier mode with a random phase,

$$\widehat{\omega}_{\mathbf{k}}(0) = \sqrt{\frac{k\,E(k, 0)}{\pi}}\;e^{\,i\zeta(\mathbf{k})}, \tag{10}$$

where the phases ζ are uniformly distributed on $[0, 2\pi)$ and satisfy $\zeta(-\mathbf{k}) = -\zeta(\mathbf{k})$, so that ω is real, and $\widehat{\omega}_{\mathbf{0}} = 0$.

Equation (6) is solved with a dealiased Fourier pseudo-spectral method on a 2048×2048 grid, using a third-order Runge–Kutta scheme for the nonlinear term and the Crank–Nicolson scheme for the viscous term. Simulations are performed at $Re = 1000$, 4000, and 8000 with time steps $\Delta t = 5 \times 10^{-4}$ ($Re \leq 4000$) and $\Delta t = 2.5 \times 10^{-4}$ ($Re = 8000$), from the same initial field, up to $T = 4$; snapshots are stored every $\Delta t_s = 0.01$. The high-resolution targets are obtained by sharp spectral truncation to the 128×128 grid,

$$\widehat{\omega^{128}}_{\mathbf{k}} = \begin{cases} \widehat{\omega^{2048}}_{\mathbf{k}}, & -64 \leq k_x, k_y < 64, \\ 0, & \text{otherwise}, \end{cases} \tag{11}$$

which retains exactly the Fourier modes representable on the target grid without aliasing.

3.1.3 Forced versus freely decaying turbulence

The two configurations pose different reconstruction problems, summarised in Table 1. In Kolmogorov flow, the forcing of Eq. (4) continuously injects energy at $k_f = 4$, which equals the nominal Nyquist wavenumber $k_c = 4$ of the 8×8 observations. A large part of the vorticity variance is therefore carried by scales that the coarse grid can at least partially represent, and the flow is statistically stationary, so that training and test snapshots share the same statistics. In DHIT there is no forcing and the energy-containing scales are initially concentrated around $k_0 = 10$, well above k_c. The coarse observations then sample a field dominated by unresolved vortices, and the statistics evolve in time as the spectral peak moves to larger scales and the vortices merge and decay. A reconstruction method must therefore infer most of the vortical content rather than refine it, and must do so across a continuously changing flow state.

3.2 Spectral Representation and Energy Statistics

Because the computational domain $\Omega = (0, 2\pi)^2$ is doubly periodic, the vorticity field admits a Fourier series representation

$$\omega(x, y, t) = \sum_{(k_x, k_y) \in \mathbb{Z}^2} \widehat{\omega}_{k_x, k_y}(t)\, e^{i(k_x x + k_y y)}, \tag{12}$$

where the Fourier coefficients are given by

$$\widehat{\omega}_{k_x, k_y}(t) = \frac{1}{(2\pi)^2} \int_\Omega \omega(x, y, t)\, e^{-i(k_x x + k_y y)}\, dx\, dy. \tag{13}$$

For a field sampled on an $n \times n$ grid, these coefficients are evaluated for $-n/2 \leq k_x, k_y < n/2$ by the discrete Fourier transform divided by n^2. With this normalisation, Parseval's identity reads $\langle \omega^2 \rangle = \sum_{\mathbf{k}} |\widehat{\omega}_{\mathbf{k}}|^2$. The isotropic energy spectrum $E(k)$ is defined by summing the modal kinetic energy over shells of unit width, $k - \frac{1}{2} \leq \|\mathbf{k}\|_2 < k + \frac{1}{2}$, of the radial wavenumber $\|\mathbf{k}\|_2 = \sqrt{k_x^2 + k_y^2}$,

$$E(k) = \frac{1}{2} \sum_{\substack{(k_x,k_y)\in\mathbb{Z}^2 \\ \|\mathbf{k}\|_2 \approx k}} |\widehat{\mathbf{u}}_{\mathbf{k}}|^2 , \tag{14}$$

where $\widehat{\mathbf{u}}_{\mathbf{k}}$ denotes the Fourier coefficient of the velocity field at wavenumber $\mathbf{k} = (k_x, k_y)$, so that $\sum_k E(k) = K$ with K the kinetic energy of Eq. (7). The dual-cascade theory of two-dimensional turbulence [31, 40] predicts that energy is transferred toward large scales (small k), whereas the enstrophy Z of Eq. (7) is transferred toward small scales (large k). The enstrophy spectrum is correspondingly defined as

$$\mathcal{E}(k) = \frac{1}{2} \sum_{\substack{(k_x,k_y)\in\mathbb{Z}^2 \\ \|\mathbf{k}\|_2 \approx k}} |\widehat{\omega}_{\mathbf{k}}|^2 = \frac{1}{2} \sum_{\substack{(k_x,k_y)\in\mathbb{Z}^2 \\ \|\mathbf{k}\|_2 \approx k}} \|\mathbf{k}\|_2^2 \, |\widehat{\mathbf{u}}_{\mathbf{k}}|^2 \approx k^2 E(k), \tag{15}$$

where the second equality uses $|\widehat{\omega}_{\mathbf{k}}|^2 = \|\mathbf{k}\|_2^2 |\widehat{\mathbf{u}}_{\mathbf{k}}|^2$, which holds for a divergence-free velocity field, and the approximation replaces $\|\mathbf{k}\|_2$ by the shell wavenumber k. These spectra quantify the distribution of kinetic energy and enstrophy across spatial scales, which pointwise error metrics such as the MSE do not resolve, and serve as the primary diagnostics of the physical fidelity of super-resolved fields.

3.3 Dataset and Discretization

For Kolmogorov-forced turbulence, the dataset consists of $N = 1001$ temporally separated snapshots of the statistically stationary vorticity field, generated by direct numerical simulation of the forced Navier–Stokes system (1)–(4) on the periodic domain Ω. Each snapshot is discretized on a uniform Cartesian grid with $n_h = 128$ points per spatial dimension, yielding a high-resolution vorticity sample

$$\omega^{128} \in \mathbb{R}^{128\times 128}. \tag{16}$$

The simulation is advanced until a statistically stationary regime is reached, as assessed through the temporal behavior of the energy spectrum and enstrophy. Following this initial transient, temporally separated snapshots are collected. Although these snapshots are obtained at different times, they originate from the same numerical simulation and are not assumed to constitute statistically independent realizations. The resulting dataset contains large-scale coherent vortical structures driven by the forcing and fine-scale turbulent fluctuations generated by nonlinear advection, spanning a broad range of resolved spatial wavenumbers. The snapshots are partitioned at random into 700 training, 100 validation, and 201 test samples.

For decaying homogeneous isotropic turbulence, the dataset at each Reynolds number, $Re \in \{1000, 4000, 8000\}$, consists of $N = 401$ snapshots at $t = 0, 0.01, \ldots, 4$ from one realisation of the simulation described in Section 3.1.2, each truncated to 128×128 according to Eq. (11). The vorticity is normalised by the standard deviation of the training fields of the respective dataset. The snapshots are partitioned into 280 training, 40 validation, and 81 test samples, following the same 70/10/20 protocol as for the Kolmogorov flow; the same snapshot indices are used at all three Reynolds numbers, and a separate instance of each model is trained for each Reynolds number. Because

all snapshots of a dataset belong to a single trajectory sampled at $\Delta t_s = 0.01$, test snapshots generally have temporally adjacent training snapshots; the DHIT results of Section 6.2 therefore measure reconstruction within a trajectory spanning early, intermediate, and late stages of the decay.

Table 1: Simulation and dataset parameters of the two turbulence configurations.

	Kolmogorov-forced	DHIT
Forcing	$\sin(k_f y)\,\mathbf{e}_x$, $k_f = 4$	none ($f \equiv 0$)
Statistical state	stationary	freely decaying
Reynolds number	$Re = 1/\nu = 40$	$Re = 1/\nu \in \{1000, 4000, 8000\}$
Initial condition	developed stationary state	Eq. (9), $k_0 = 10$
Simulation grid	128^2	2048^2, truncated to 128^2
Snapshots N	1001	401 per Re, $t \in [0, 4]$
Train / validation / test	700/100/201	280/40/81 per Re

Low-resolution observations are constructed for both datasets with the same restriction from 128×128 to 8×8, defined in Section 3.4. For both datasets, the training set is used exclusively for model parameter optimization, while the validation set is employed for model selection and checkpoint selection during training. The test set remains held out throughout training and model selection and is used only for the final performance evaluation. Identical data partitions and preprocessing procedures are maintained across SIMR-NO and all baseline models to ensure a consistent and fair comparison.

3.4 Measurement Model and Inverse Problem Formulation

In the setting considered here, the high-resolution field ω^{128} is not directly observed. Instead, only a spatially coarse observation is available, obtained by applying a spatial restriction operator

$$a^{r_c} = \mathcal{M}_{128 \to r_c}(\omega^{128}) \in \mathbb{R}^{r_c \times r_c}, \tag{17}$$

where $\mathcal{M}_{128 \to r_c}$ maps the 128×128 grid to a coarser $r_c \times r_c$ grid. In this work, we consider $r_c = 8$, corresponding to a $16\times$ linear downsampling factor, so that

$$a^8 \in \mathbb{R}^{8 \times 8}. \tag{18}$$

This setting retains only 64 spatial samples from the original 16,384 grid points, that is, 0.39% of the high-resolution samples. The present setting remains challenging because fine-scale flow structures are not directly resolved by the available measurements. Specifically, the restriction operator evaluates the cubic B-spline interpolant $S_3[\omega^{128}]$ of the grid values at the coarse sample positions,

$$a^8_{i,j} = \big[\mathcal{M}_{128 \to 8}(\omega^{128})\big]_{i,j} = S_3[\omega^{128}](\xi_i, \xi_j), \qquad \xi_i = \frac{127\, i}{7}, \qquad i, j = 0, \ldots, 7, \tag{19}$$

where positions are measured in units of the fine-grid spacing, with the fine-grid points at $0, 1, \ldots, 127$. This is the spline resampling of `scipy.ndimage.zoom` with order 3, which is used for both datasets. The coarse samples are spaced by $127/7 \approx 18.1$ fine-grid cells, and the first and last samples coincide with the first and last fine-grid points, so the sampling does not exploit the periodicity of Ω. Like pointwise subsampling, Eq. (19) applies no low-pass filter before sampling: unlike ideal spectral truncation, it does not retain exclusively the low-wavenumber Fourier modes, and unresolved high-wavenumber components are folded into the resolved range of the coarse grid. Consequently, the coarse observation cannot be interpreted as a low-pass representation of the vorticity field. The nominal Nyquist wavenumber of the coarse grid is $k_c = r_c/2 = 4$ in each direction, whereas the reference field contains structures at wavenumbers up to 64 that cannot be uniquely determined from the coarse measurements.

The resulting inverse reconstruction problem,

$$\text{find } \omega^{128} \text{ such that } \mathcal{M}_{128\to 8}(\omega^{128}) = a^8, \tag{20}$$

is ill-posed in the sense of Hadamard [6]: because $\mathcal{M}_{128\to 8}$ is a linear map from $\mathbb{R}^{128\times 128}$ to $\mathbb{R}^{8\times 8}$, its null space has dimension at least 16,320, and infinitely many high-resolution fields are consistent with the same coarse observation. A unique reconstruction therefore requires prior information, which in the present work is supplied by the training data.

To provide an initial approximation for the learned reconstruction, the coarse observation is first upsampled to the target resolution by bicubic interpolation,

$$\widetilde{\omega}^{128} = \mathcal{U}_{8\to 128}(a^8) \in \mathbb{R}^{128\times 128}, \tag{21}$$

where $\mathcal{U}_{r\to r'}$ denotes bicubic convolution interpolation from an $r \times r$ to an $r' \times r'$ grid, as implemented in PyTorch [81] with the corner-aligned convention (`align_corners=True`), which places coarse sample i at fine-grid position $127i/7$, the position at which Eq. (19) takes it, so that $\widetilde{\omega}^{128}$ passes through the coarse samples. The pseudo-high-resolution field is therefore a smooth estimate whose large-scale structures are distorted by aliasing. In general it does not reproduce the low-wavenumber Fourier coefficients of ω^{128} and contains no information about the unresolved fine-scale structures.

Accordingly, the reconstruction error of the interpolated field is

$$r^{128} \coloneqq \omega^{128} - \widetilde{\omega}^{128} = \sum_{\mathbf{k}\in\mathbb{K}_{128}} \left(\widehat{\omega}_{\mathbf{k}} - \widehat{\widetilde{\omega}}_{\mathbf{k}}\right) e^{i(k_x x + k_y y)}, \tag{22}$$

where $\mathbb{K}_{128} = \{(k_x, k_y) \in \mathbb{Z}^2 : -64 \le k_x, k_y < 64\}$ is the set of wavenumbers resolved on the high-resolution grid, and $\widehat{\omega}_{\mathbf{k}}$ and $\widehat{\widetilde{\omega}}_{\mathbf{k}}$ are the Fourier coefficients of the reference and interpolated fields. Because of aliasing and the interpolation kernel, r^{128} contains contributions at both low and high wavenumbers. The learned models must therefore correct the large-scale errors of $\widetilde{\omega}^{128}$ as well as infer the unresolved fine-scale structures.

3.5 Super-Resolution as Operator Learning

The turbulence super-resolution problem is formalized as learning an operator $\mathcal{G}_\theta$, parameterized by neural network weights $\theta \in \Theta$, that maps the pseudo-high-resolution estimate to an approximation of the high-resolution vorticity field,

$$\mathcal{G}_\theta : \widetilde{\omega}^{128} \longmapsto \widehat{\omega}^{128} \approx \omega^{128}, \tag{23}$$

where $\widehat{\omega}^{128} \coloneqq \mathcal{G}_\theta(\widetilde{\omega}^{128})$ denotes the reconstructed field. (The convolutional baselines of Section 5 take a^8 instead of $\widetilde{\omega}^{128}$ as input; both inputs are determined by the same coarse observation, and no model has access to any other information about ω^{128}.) The operator is trained by minimizing the empirical risk

$$\mathcal{L}(\theta) = \frac{1}{N_{\text{train}}} \sum_{i=1}^{N_{\text{train}}} \text{MSE}\big(\mathcal{G}_\theta\big(\widetilde{\omega}_i^{128}\big)\,,\, \omega_i^{128}\big)\,, \tag{24}$$

where the mean squared error over the 128×128 grid points is defined in Eq. (42) and equals the squared $L^2(\Omega)$ norm of the error divided by $|\Omega| = (2\pi)^2$, up to the quadrature error of the grid. Because $\mathcal{M}_{128\to 8}$ is not injective, ω^{128} is not in general a function of $\widetilde{\omega}^{128}$, and no operator can reproduce every high-resolution field from its coarse observation. Over the data distribution, the minimizer of the expected squared error is the conditional mean $\mathbb{E}[\omega^{128} \mid \widetilde{\omega}^{128}]$, which is therefore the target approximated by $\mathcal{G}_\theta$. The learning problem (23)–(24) differs from standard image super-resolution in a physically important respect: a useful reconstruction must not only achieve a low pointwise error but also reproduce the spectral content of ω^{128}, including the energy and enstrophy spectra $E(k)$ and $\mathcal{E}(k)$ of Eqs. (14)–(15). The design of SIMR-NO in Section 4 therefore combines spectral and multiscale architectural elements. Reconstruction quality is assessed with the pointwise, structural, spectral, and localized metrics defined in Section 5.4, evaluated on the held-out test set.

4 Mathematical Formulation of SIMR-NO

This section presents the mathematical formulation of the spectrally-informed multi-resolution neural operator (SIMR-NO). We first state the operator-learning framework and the residual formulation, then the hierarchical multi-resolution decomposition, the stage architecture with its local refinement modules, the spectrally gated convolution, and the training procedure. All notation follows Section 3, and the complete forward pass is summarized in Algorithm 1. An overview of the full architecture is provided in Figure 1.

4.1 Operator Learning Framework

Let $\mathcal{H}_r \coloneqq \mathbb{R}^{r\times r}$ denote the space of vorticity fields sampled on a uniform $r \times r$ Cartesian grid over the periodic domain $\Omega = (0, 2\pi)^2$, equipped with the discrete L^2 norm. Following Section 3, the target of the reconstruction is the conditional mean $\mathcal{G}^\star(\widetilde{\omega}^{128}) = \mathbb{E}[\omega^{128} \mid \widetilde{\omega}^{128}]$, where $\widetilde{\omega}^{128} = \mathcal{U}_{8\to 128}(a^8)$ is the bicubic pseudo-high-resolution field of Eq. (21). The operator $\mathcal{G}^\star : \mathcal{H}_{128} \to \mathcal{H}_{128}$ is unknown and is approximated by a parametric family $\{\mathcal{G}_\theta\}_{\theta\in\Theta}$ of neural operators, with $\Theta \subset \mathbb{R}^{n_\theta}$.

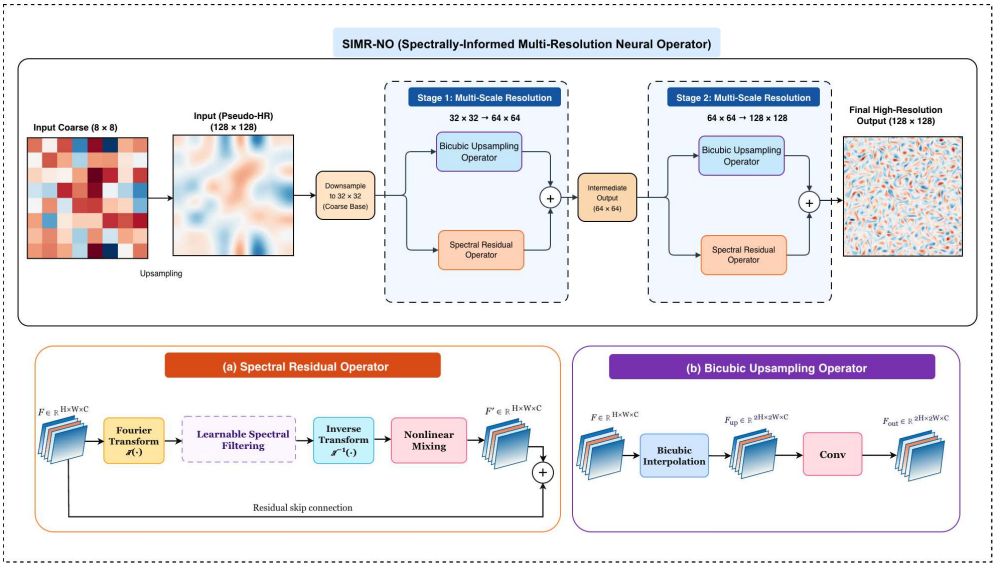


Fig. 1: SIMR-NO architecture. The pseudo-high-resolution field $\widetilde{\omega}^{128}$ is downsampled to a^{32} (Eq. (28)) and reconstructed in two stages, $32 \to 64$ and $64 \to 128$ (Eqs. (30)–(31)), each adding a learned spectral residual, scaled by α_s, to a bicubic upsampling of its input. (a) Spectral residual operator; (b) bicubic upsampling operator.

SIMR-NO approximates $\mathcal{G}^\star$ in residual form. With the residual $r^{128} = \omega^{128} - \widetilde{\omega}^{128}$ of Eq. (22),

$$\mathcal{G}^\star(\widetilde{\omega}^{128}) = \widetilde{\omega}^{128} + \mathcal{R}^\star(\widetilde{\omega}^{128}), \qquad \mathcal{R}^\star(\widetilde{\omega}^{128}) = \mathbb{E}[r^{128} \mid \widetilde{\omega}^{128}], \tag{25}$$

where $\mathcal{R}^\star$ is the residual correction operator to be approximated by the network. By Parseval's identity, the residual energy satisfies

$$\left\| r^{128} \right\|^2_{L^2(\Omega)} = (2\pi)^2 \sum_{\mathbf{k} \in \mathbb{K}_{128}} \left| \widehat{\omega}_{\mathbf{k}} - \widehat{\widetilde{\omega}}_{\mathbf{k}} \right|^2, \tag{26}$$

with the Fourier coefficients of Eq. (13). Because the coarse observations are aliased (Section 3), the residual energy is not confined to wavenumbers above the coarse-grid Nyquist limit, and its norm is not necessarily smaller than that of ω^{128}: for DHIT, the mean relative ℓ_2 error of the bicubic field exceeds one (Table 4). The residual form is therefore adopted not because the residual is small, but because it provides an initialization at the interpolated field and lets the network represent corrections to it rather than the entire field. Such residual corrections are applied at each refinement stage of SIMR-NO.

4.2 Hierarchical Multi-Resolution Decomposition and Residual Correction

The central architectural principle of SIMR-NO is to factorize the reconstruction operator into a cascade of stage-wise operators, each of which doubles the spatial resolution. Let $\{r_0, r_1, r_2\} = \{32, 64, 128\}$ denote the sequence of intermediate resolutions. The two-stage reconstruction is

$$\mathcal{G}^{\mathrm{SIMR}} = \mathcal{G}^{(2)}_{\theta_2} \circ \mathcal{G}^{(1)}_{\theta_1} \circ \mathcal{U}_{128\to 32}, \tag{27}$$

followed by the refinement head of Eq. (32), where $\mathcal{G}^{(s)}_{\theta_s} : \mathcal{H}_{r_{s-1}} \to \mathcal{H}_{r_s}$ is the learned stage-s operator with parameters θ_s. The 32×32 input to Stage 1 is obtained from the pseudo-high-resolution field by bicubic downsampling,

$$a^{32} = \mathcal{U}_{128\to 32}(\widetilde{\omega}^{128}) \in \mathbb{R}^{32\times 32}. \tag{28}$$

Since a^{32} is computed from $\widetilde{\omega}^{128}$, it carries no information beyond that contained in the 64 coarse samples; the hierarchy distributes the reconstruction over several resolutions but does not add observational information. On an $r \times r$ grid, wavenumbers up to $r/2$ are representable, so each stage extends the representable range by one octave, from $r_{s-1}/2$ to $r_s/2$, and the upscaling factor that each learned stage must bridge is 2 rather than the 16 of a single-stage mapping from the coarse observation. By analogy with multiresolution analysis [82] and multigrid methods [37], this decomposition is expected to make each stage-wise mapping easier to learn than a direct mapping across the full resolution range; the present work assesses this expectation empirically rather than proving it.

Each stage $s \in \{1, 2\}$ is formulated as a residual correction. Given the input field $a^{r_{s-1}} \in \mathcal{H}_{r_{s-1}}$, a smooth base estimate at the target resolution is first obtained by bicubic interpolation, $b^{r_s} = \mathcal{U}_{r_{s-1}\to r_s}(a^{r_{s-1}})$. The stage output is then

$$\widehat{\omega}^{r_s} = b^{r_s} + \alpha_s \, \mathcal{R}^{(s)}_{\theta_s}(a^{r_{s-1}}), \tag{29}$$

where $\alpha_s \in \mathbb{R}$ is a learnable scalar, initialized to 1, that controls the magnitude of the residual correction, and $\mathcal{R}^{(s)}_{\theta_s} : \mathcal{H}_{r_{s-1}} \to \mathcal{H}_{r_s}$ is the learned spectrally informed residual operator. Whenever the residual operator is small, the stage output reduces to the bicubic interpolation b^{r_s}, which provides a well-defined starting point for optimization. Applying (29) to both stages, the two-stage reconstruction is

$$\widehat{\omega}^{64} = \mathcal{U}_{32\to 64}(a^{32}) + \alpha_1 \, \mathcal{R}^{(1)}_{\theta_1}(a^{32}), \tag{30}$$

$$\widehat{\omega}^{128} = \mathcal{U}_{64\to 128}(\widehat{\omega}^{64}) + \alpha_2 \, \mathcal{R}^{(2)}_{\theta_2}(\widehat{\omega}^{64}). \tag{31}$$

After Stage 2, a lightweight refinement head $\mathcal{C}_{\theta_3}$ applies a residual correction to the 128×128 output,

$$\widehat{\omega}^{128}_{\mathrm{final}} = \widehat{\omega}^{128} + \mathcal{C}_{\theta_3}(\widehat{\omega}^{128}), \tag{32}$$

where $\mathcal{C}_{\theta_3}$ is the three-layer convolutional network $\mathrm{Conv}_{3\times3}^{d_h\to1} \circ \mathrm{GELU} \circ \mathrm{Conv}_{3\times3}^{d_h\to d_h} \circ \mathrm{GELU} \circ \mathrm{Conv}_{3\times3}^{1\to d_h}$ with hidden width $d_h = 32$. Its receptive field of 7×7 grid points restricts it to local corrections, which complement the global spectral operators of the stages.

4.3 Stage Architecture: Lifting, Spectral Blocks, and Local Refinement

Each stage s is implemented as a `SIMRStage` module and processes its input through four sequential components. First, the input field $a^{r_{s-1}}$ is bicubic-interpolated to the target resolution to form the base b^{r_s}, which is augmented with a fixed positional encoding $\boldsymbol{\phi}$ consisting of the Cartesian coordinates of the grid points, mapped linearly to $x, y \in [-1, 1]$, and Fourier features,

$$\boldsymbol{\phi}(x,y) = \big[x,\ y,\ \{\sin(\pi i x), \cos(\pi i x), \sin(\pi i y), \cos(\pi i y)\}_{i=1}^{n_f}\big]^\top \in \mathbb{R}^{c_{\mathrm{ff}}}, \tag{33}$$

with $n_f = 6$ Fourier frequencies giving $c_{\mathrm{ff}} = 2 + 4n_f = 26$ positional channels. The augmented pointwise input $[b^{r_s}(x,y),\ \boldsymbol{\phi}(x,y)]^\top \in \mathbb{R}^{1+c_{\mathrm{ff}}}$ is lifted to the working channel width d_s by a pointwise linear layer $\mathcal{P}_s$, yielding the initial feature field $\mathbf{v}^{(0)} = \mathcal{P}_s\,[b^{r_s}, \boldsymbol{\phi}]^\top \in \mathcal{H}_{r_s}^{d_s}$. Stage 1 uses width $d_1 = 64$ with $n_1 = 3$ FNO blocks and $l_1 = 2$ local refinement blocks; Stage 2 uses width $d_2 = 80$ with $n_2 = 5$ FNO blocks and $l_2 = 4$ local refinement blocks, giving greater capacity to the stage that operates at the target resolution. Second, the lifted features are processed through n_s sequentially composed FNO blocks, each applying a spectrally gated convolution in parallel with a local 1×1 bypass:

$$\mathbf{v}^{(\ell+1)} = \mathrm{GELU}\Big(\mathcal{K}_{\mathrm{gate}}^{(\ell)}[\mathbf{v}^{(\ell)}] + W^{(\ell)}\mathbf{v}^{(\ell)}\Big), \qquad \ell = 0, \dots, n_s - 1, \tag{34}$$

where $W^{(\ell)} \in \mathbb{R}^{d_s \times d_s}$ is a pointwise 1×1 convolution and $\mathcal{K}_{\mathrm{gate}}^{(\ell)}$ is the spectrally gated operator defined below. Third, after the FNO blocks, the feature field passes through l_s residual local refinement blocks, each consisting of three zero-padded 3×3 convolutions with GELU activations and a skip connection,

$$\mathbf{v} \leftarrow \mathbf{v} + \mathrm{Conv}_{3\times3} \circ \mathrm{GELU} \circ \mathrm{Conv}_{3\times3} \circ \mathrm{GELU} \circ \mathrm{Conv}_{3\times3}(\mathbf{v}), \tag{35}$$

which can represent spatially localized features that are not contained in the truncated Fourier basis of the spectral layers. Fourth, the refined feature field is projected back to the scalar vorticity space through two pointwise linear layers, $\widehat{r}^{r_s} = \mathcal{Q}_2\,\mathrm{GELU}(\mathcal{Q}_1\,\mathbf{v}) \in \mathcal{H}_{r_s}$, where $\mathcal{Q}_1 \in \mathbb{R}^{128\times d_s}$ and $\mathcal{Q}_2 \in \mathbb{R}^{1\times128}$, giving the stage residual $\widehat{r}^{r_s}$ used in (29).

4.4 Spectrally Gated Convolution

The operator $\mathcal{K}_{\mathrm{gate}}^{(\ell)}$ in (34) is the main spectral component of SIMR-NO. Let $\widehat{\mathbf{v}}(\mathbf{k}) \in \mathbb{C}^{d_s}$ denote the discrete Fourier transform of the feature field at wavenumber $\mathbf{k} = (k_x, k_y)$, computed with the orthonormal real-to-complex transform `rfft2`, which stores the coefficients with $k_x \geq 0$ along its half-spectrum axis. A spectral convolution with m retained modes per direction computes

$$\widehat{\mathbf{y}}(\mathbf{k}) = W_\theta(\mathbf{k})\,\widehat{\mathbf{v}}(\mathbf{k}), \qquad \mathbf{k} \in \mathbb{K}_m \coloneqq \{(k_x, k_y) : 0 \le k_x < m,\ 0 \le k_y < m\}, \tag{36}$$

where $W_\theta(\mathbf{k}) \in \mathbb{C}^{d_s \times d_s}$ is a learned complex Fourier multiplier, and $\widehat{\mathbf{y}}(\mathbf{k}) = \mathbf{0}$ for all other coefficients stored by `rfft2`. Because the output is real, the conjugate-symmetric coefficients $\widehat{\mathbf{y}}(-\mathbf{k}) = \overline{\widehat{\mathbf{y}}(\mathbf{k})}$ are implied, so the layer acts on the wavenumbers with $k_x k_y \ge 0$ and $|k_x|, |k_y| < m$; this is the retained-mode set of the FNO implementation that is also used for the FNO baseline. SIMR-NO augments (36) with a scalar radial gating function $g : [0, 1) \to (0, 1)$ applied to the normalized radial wavenumber

$$\rho(\mathbf{k}) \coloneqq \frac{\|\mathbf{k}\|_2}{\max_{\mathbf{k}' \in \mathbb{K}_m} \|\mathbf{k}'\|_2 + \varepsilon} \in [0, 1), \qquad \varepsilon = 10^{-12}, \tag{37}$$

where $\max_{\mathbf{k}' \in \mathbb{K}_m} \|\mathbf{k}'\|_2 = \sqrt{2}\,(m-1)$. The gate is parameterized as a two-layer perceptron with hidden dimension $d_g = 32$,

$$g(\rho) = \mathrm{Sigmoid}(A_2\,\mathrm{GELU}(A_1\rho + b_1) + b_2)\,, \tag{38}$$

where $A_1 \in \mathbb{R}^{d_g \times 1}$, $b_1 \in \mathbb{R}^{d_g}$, $A_2 \in \mathbb{R}^{1 \times d_g}$, and $b_2 \in \mathbb{R}$ are learnable parameters. The sigmoid constrains $g(\rho) \in (0, 1)$, so that no retained mode is completely suppressed. The spectrally gated convolution computes

$$\widehat{\mathbf{y}}(\mathbf{k}) = g(\rho(\mathbf{k}))\,W_\theta(\mathbf{k})\,\widehat{\mathbf{v}}(\mathbf{k}), \qquad \mathbf{k} \in \mathbb{K}_m, \tag{39}$$

with the spatial output $\mathbf{y} = \mathcal{F}^{-1}[\widehat{\mathbf{y}}]$ recovered via `irfft2`. All Fourier operations are performed in single precision with automatic mixed precision disabled. Stage 1 uses $m^{(1)} = 16$ modes and Stage 2 uses $m^{(2)} = 32$ modes.

Since g depends on $\mathbf{k}$ only through $\|\mathbf{k}\|_2$, it introduces one shared, learned weighting of all retained modes by their scale, in addition to the mode-wise multipliers $W_\theta(\mathbf{k})$; the multipliers remain free to represent directional structure, such as the anisotropy imposed by the Kolmogorov forcing of Eq. (4). Because g is learned, the relative weight given to spectral corrections at different scales is determined by the training data rather than prescribed, without explicit supervision on the energy spectrum.

4.5 Training Objective and Optimization

The complete SIMR-NO model, comprising Stage 1 parameters θ_1, Stage 2 parameters θ_2, final refinement head parameters θ_3, and learnable scaling parameters α_1, α_2, is trained end-to-end by minimizing the loss of Eq. (24) with $\mathcal{G}_\theta(\widetilde{\omega}^{128}) = \widehat{\omega}^{128}_{\mathrm{final}}$, that is, the mean squared error of Eq. (42) between the final reconstruction and the reference field, averaged over the training set (over each mini-batch during optimization). The loss is applied only to the final output, with gradients propagating through (32), (31), and (30) simultaneously, so that the intermediate field $\widehat{\omega}^{64}$ is not constrained to match any particular 64×64 target. The model is trained for 300 epochs with a batch size of 16, and the checkpoint with the lowest validation MSE is used for evaluation. The test set is not used during training or model selection. The complete procedure is summarized in Algorithm 1.

Algorithm 1 SIMR-NO: Forward Pass, Training, and Inference

Require: Training, validation, and test sets $\mathcal{D}_{\text{train}}$, $\mathcal{D}_{\text{val}}$, $\mathcal{D}_{\text{test}}$ of pairs $(\widetilde{\omega}^{128}, \omega^{128})$; number of epochs E

Ensure: Trained parameters $(\theta_1^\star, \theta_2^\star, \theta_3^\star, \alpha_1^\star, \alpha_2^\star)$

Initialize Stage 1 $\mathcal{G}_{\theta_1}^{(1)}$, Stage 2 $\mathcal{G}_{\theta_2}^{(2)}$, and refinement head $\mathcal{C}_{\theta_3}$
Initialize the optimizer and learning-rate scheduler; BestVal $\leftarrow +\infty$
for epoch $e = 1$ **to** E **do**
  for each mini-batch $\mathcal{B} \subset \mathcal{D}_{\text{train}}$ **do**
    // Downsample pseudo HR input to 32×32
    $a^{32} \leftarrow \mathcal{U}_{128\to 32}(\widetilde{\omega}^{128})$
    // Stage 1: 32→64
    $\widehat{\omega}^{64} \leftarrow \mathcal{U}_{32\to 64}(a^{32}) + \alpha_1 \, \mathcal{R}_{\theta_1}^{(1)}(a^{32})$
    // Stage 2: 64→128
    $\widehat{\omega}^{128} \leftarrow \mathcal{U}_{64\to 128}(\widehat{\omega}^{64}) + \alpha_2 \, \mathcal{R}_{\theta_2}^{(2)}(\widehat{\omega}^{64})$
    // Final pixel-level refinement
    $\widehat{\omega}_{\text{final}}^{128} \leftarrow \widehat{\omega}^{128} + \mathcal{C}_{\theta_3}(\widehat{\omega}^{128})$
    // Compute training loss and update
    $\mathcal{L}_{\mathcal{B}} \leftarrow \frac{1}{|\mathcal{B}|} \sum_{i\in\mathcal{B}} \text{MSE}\big(\widehat{\omega}_{\text{final},i}^{128}, \omega_i^{128}\big)$
    Backpropagate $\mathcal{L}_{\mathcal{B}}$ and update the parameters
  end for
  Step the learning-rate scheduler
  Evaluate the validation loss $\mathcal{L}_{\text{val}}$ on $\mathcal{D}_{\text{val}}$
  if $\mathcal{L}_{\text{val}} <$ BestVal **then**
    BestVal $\leftarrow \mathcal{L}_{\text{val}}$; save the current parameters
  end if
end for
Restore the parameters with the lowest validation loss
Evaluate the selected model on $\mathcal{D}_{\text{test}}$

Inference: Given a new observation $\widetilde{\omega}^{128}$, apply the two-stage reconstruction and the final refinement to obtain $\widehat{\omega}_{\text{final}}^{128}$.

4.6 Design Rationale and Error Propagation

The SIMR-NO architecture is motivated by two principles, which we state as design rationale rather than as guarantees of accuracy. First, the residual decomposition (25) initializes each stage at a deterministic interpolant, so that the network learns corrections to an explicit estimate; the learnable scale α_s controls the size of these corrections. As discussed after Eq. (26), the residual is not necessarily small in the present setting, so this formulation does not by itself reduce the complexity of the learning target. Second, the hierarchical factorization (27) follows the coarse-to-fine organization of multiresolution analysis [82, 83], with each stage doubling the resolution.

The factorization also determines how errors of the individual stages combine. Let $S_1(a^{32}) = \widehat{\omega}^{64}$, $S_2(\widehat{\omega}^{64}) = \widehat{\omega}^{128}$, and $S_3(\widehat{\omega}^{128}) = \widehat{\omega}_{\text{final}}^{128}$ denote the maps of Eqs. (30),

(31), and (32), and let $\omega^{64}_{\text{ref}} \in \mathcal{H}_{64}$ be any reference field at the intermediate resolution, for example a restriction of ω^{128}. Each map is a composition of affine maps, interpolation operators, Fourier transforms, and GELU and sigmoid functions, all of which are Lipschitz continuous on finite-dimensional spaces; S_2 and S_3 therefore have finite Lipschitz constants L_2 and L_3. Defining the stage errors

$$\epsilon_1 = \left\| S_1(a^{32}) - \omega^{64}_{\text{ref}} \right\|, \qquad \epsilon_2 = \left\| S_2(\omega^{64}_{\text{ref}}) - \omega^{128} \right\|, \qquad \epsilon_3 = \left\| S_3(\omega^{128}) - \omega^{128} \right\|, \tag{40}$$

the triangle inequality gives

$$\left\| \widehat{\omega}^{128}_{\text{final}} - \omega^{128} \right\| \leq L_3 L_2 \, \epsilon_1 + L_3 \, \epsilon_2 + \epsilon_3, \tag{41}$$

in the discrete L^2 norm. The bound shows that an error of Stage 1 is propagated through, and possibly amplified by, the subsequent stages, which motivates training all stages jointly on the final output rather than independently. It does not imply that the factorized operator is more accurate than a single-stage operator; this comparison is assessed empirically in Section 6.

5 Experimental Setup

This section describes the setup used to compare SIMR-NO with its baselines. All experiments are conducted independently on the Kolmogorov-forced and the decaying homogeneous isotropic turbulence datasets of Section 3: for each dataset, every model is trained, selected, and tested on the same partitions. The same training protocol is applied to all models, so that performance differences can be attributed to architectural design rather than to differences in the training setup.

5.1 Dataset and Preprocessing

The Kolmogorov dataset contains $N = 1001$ snapshots of the statistically stationary vorticity field, and each of the three DHIT datasets contains $N = 401$ snapshots of the decaying flow at $Re = 1000$, 4000, or 8000, all represented as $\omega^{128} \in \mathbb{R}^{128\times128}$ on the periodic domain $\Omega = (0, 2\pi)^2$ (Section 3.3, Table 1). Each high-resolution snapshot is restricted to the coarse observation a^8 of Eq. (19), which is upsampled to the pseudo-high-resolution field $\widetilde{\omega}^{128}$ of Eq. (21); this field is also reported as the non-learned bicubic baseline.

Each dataset is partitioned into mutually exclusive training, validation, and test sets. The training set is used exclusively for parameter optimization, the validation set for checkpoint selection based on the lowest validation mean squared error (MSE), and the test set only for the final evaluation. Identical partitions and preprocessing are used for SIMR-NO and all baseline models.

5.2 Baseline Methods

To deliver an extensive evaluation that includes different evaluation methods, SIMR-NO receives comparisons against four foundational techniques, which include

traditional interpolation methods and neural operator learning and convolutional super-resolution. The non-learned baseline is the bicubic field $\widetilde{\omega}^{128}$ of Eq. (21). It adds no spectral information beyond the coarse observation and establishes the reference level that the learned methods must improve upon. The Fourier Neural Operator (FNO) [11] serves as a complete neural operator system that executes global convolution through Fourier space while utilizing its trained spectral multipliers. The state-of-the-art forward PDE modeling operator learning system uses FNO as its main neural operator benchmark. The FNO baseline uses four Fourier layers of width 64 with 32 retained Fourier modes per direction and receives the same pseudo-high-resolution input $\widetilde{\omega}^{128}$ as SIMR-NO.

EDSR develops an advanced deep residual convolutional network, which enables improved image super-resolution performance. The system eliminates batch normalization from its residual blocks because this approach boosts reconstruction accuracy according to research findings in [29]. EDSR functions as a powerful single-stage convolutional system that achieves local spatial pattern detection through its multiple stacked 3×3 residual blocks. The progressive multi-scale super-resolution network LapSRN generates its output by using a Laplacian pyramid, which performs upsampling and convolutional refinement at different resolution levels [35]. The architectural design of LapSRN serves as the closest matching base model to SIMR-NO because both methods implement a hierarchical reconstruction method, which enables direct comparison of operator learning and spectral gating through convolutional evaluation. All learned models are thus conditioned on the same coarse observation, are trained on the same data splits, and are optimized with the same loss function and optimizer type, so that performance differences are attributable to architectural design rather than training disparities.

5.3 Training Configuration and Fairness Protocol

All neural network models use PyTorch as their implementation framework [81] to train the models with the Adam optimizer [84] at a 16-batch size for 300 epochs to minimize the mean-squared-error loss of Eq. (24). The learning rate and the architectural hyperparameters are set separately for each model; those of SIMR-NO are given in Section 4 and Algorithm 1.

To ensure a fair and reproducible comparison, all models are trained and evaluated using identical training, validation, and test partitions. Each model is optimized using the training set, while the validation set is used exclusively for checkpoint selection. The test set is held out until the final evaluation, ensuring that reported performance metrics are not influenced by model-selection decisions. All methods use the same input observations, reconstruction targets, and evaluation metrics.

5.4 Evaluation Metrics

Reconstruction quality is measured for each test snapshot and summarized by the mean and standard deviation over the test set. For a reconstruction $\widehat{\omega}$ of the reference field ω on the $n \times n$ grid ($n = 128$), we use the mean squared error (MSE), the relative ℓ_2 error, the peak signal-to-noise ratio (PSNR), and the structural similarity index

(SSIM),

$$\mathrm{MSE} = \frac{1}{n^2}\sum_{i,j}\left(\widehat{\omega}_{ij} - \omega_{ij}\right)^2, \qquad \mathrm{RelL}_2 = \frac{\|\widehat{\omega} - \omega\|_F}{\|\omega\|_F}, \tag{42}$$

$$\mathrm{PSNR} = 10\log_{10}\left(\frac{\max_{ij}|\omega_{ij}|^2}{\mathrm{MSE}}\right), \qquad \mathrm{SSIM} = \left\langle \frac{(2\mu_{\hat{\omega}}\mu_{\omega} + c_1)(2\sigma_{\hat{\omega}\omega} + c_2)}{(\mu_{\hat{\omega}}^2 + \mu_{\omega}^2 + c_1)(\sigma_{\hat{\omega}}^2 + \sigma_{\omega}^2 + c_2)} \right\rangle_W, \tag{43}$$

where $\|\cdot\|_F$ is the Frobenius norm, μ, σ^2, and $\sigma_{\hat{\omega}\omega}$ are local means, variances, and covariance over windows W, and c_1, c_2 are small stabilization constants. MSE and relative ℓ_2 error measure pointwise accuracy, PSNR expresses the error relative to the peak vorticity, and SSIM measures structural agreement.

Physical fidelity is assessed with three further diagnostics. First, the energy and enstrophy spectra of Eqs. (14)–(15) of each reconstruction are compared with those of the reference (for the Kolmogorov case studies, from fields normalized by their maximum absolute value). Second, to locate where errors arise in the flow, test pixels are grouped into ten quantile bins $\mathcal{B}_b$ of the reference vorticity magnitude, from the quiescent background to the vortex cores, and the mean absolute error is evaluated in each bin, in absolute terms and relative to the local vorticity. Third, for the zoomed comparisons, the relative ℓ_2 error is evaluated on a 20×20 subdomain $\mathcal{D}$:

$$\varepsilon_b = \left\langle|\widehat{\omega} - \omega|\right\rangle_{\mathcal{B}_b}, \qquad \varepsilon_b^{\mathrm{rel}} = \frac{\left\langle|\widehat{\omega} - \omega|\right\rangle_{\mathcal{B}_b}}{\left\langle|\omega|\right\rangle_{\mathcal{B}_b}}, \tag{44}$$

$$\mathrm{RelL}_2^{\mathcal{D}} = \frac{\|\widehat{\omega} - \omega\|_{F,\mathcal{D}}}{\|\omega\|_{F,\mathcal{D}}}. \tag{45}$$

6 Results

We compare SIMR-NO with bicubic interpolation, FNO [11], EDSR [29], and LapSRN [35] on the held-out test sets of both turbulence configurations. Section 6.1 reports results for Kolmogorov-forced turbulence and Section 6.2 for decaying homogeneous isotropic turbulence at $Re = 1000$, 4000, and 8000. For each flow we first present aggregate statistics over the full test set, including, for DHIT, their dependence on the Reynolds number, then representative reconstructions with spectral diagnostics, and finally localized comparisons. Section 6.3 contrasts the two regimes, and Section 6.4 summarizes the findings.

6.1 Kolmogorov-Forced Turbulence

6.1.1 Aggregate performance

Table 2 lists the mean and standard deviation of the relative ℓ_2 error, SSIM, and PSNR over the 201 test snapshots. Bicubic interpolation of the 8×8 observations yields a relative ℓ_2 error of 79.15% and an SSIM of 0.218: the coarse grid resolves only wavenumbers up to $k_c = 4$, and the interpolant cannot restore the vortical structure at finer scales. The learned models reduce this error by a factor of 2.2 (EDSR) to 3.2

(SIMR-NO). Among the baselines, LapSRN is the most accurate (0.3045), followed by FNO (0.3194) and EDSR (0.3623). SIMR-NO attains a mean relative ℓ_2 error of 0.2506 ± 0.0983, which is 17.7%, 21.5%, and 30.8% lower than that of LapSRN, FNO, and EDSR, respectively. It also attains the highest SSIM (0.8075 versus 0.7457 for LapSRN) and PSNR (24.77 dB, that is, 1.81 dB above LapSRN and 2.28 dB above FNO), and the smallest standard deviation of the relative error among the learned models.

Table 2: Kolmogorov-forced turbulence: test-set metrics (mean $\pm$ standard deviation). Arrows indicate whether lower or higher values are better; the best value in each column is shown in bold.

Method	Rel. $\ell_2 \downarrow$	SSIM $\uparrow$	PSNR (dB) $\uparrow$
Bicubic	0.7915 ± 0.0292	0.2184 ± 0.0367	14.22 ± 1.64
FNO	0.3194 ± 0.1062	0.6967 ± 0.1322	22.49 ± 3.24
EDSR	0.3623 ± 0.1054	0.6433 ± 0.1166	21.32 ± 3.04
LapSRN	0.3045 ± 0.1076	0.7457 ± 0.1207	22.96 ± 3.37
SIMR-NO	$\mathbf{0.2506 \pm 0.0983}$	$\mathbf{0.8075 \pm 0.1084}$	$\mathbf{24.77 \pm 3.68}$

Figure 2 shows the test-set distributions of the three metrics for the learned models. The improvement of SIMR-NO is a shift of the whole distribution rather than the effect of a few samples: the median relative error decreases from approximately 0.34 (EDSR), 0.30 (FNO), and 0.28 (LapSRN) to 0.23 (SIMR-NO); the medians of SSIM and PSNR increase in the same order; and the largest relative error decreases from approximately 0.75 (EDSR), 0.71 (FNO), and 0.67 (LapSRN) to 0.62 for SIMR-NO. The spread of the distributions, which reflects the variability of the flow state across test snapshots, is comparable for all learned models.

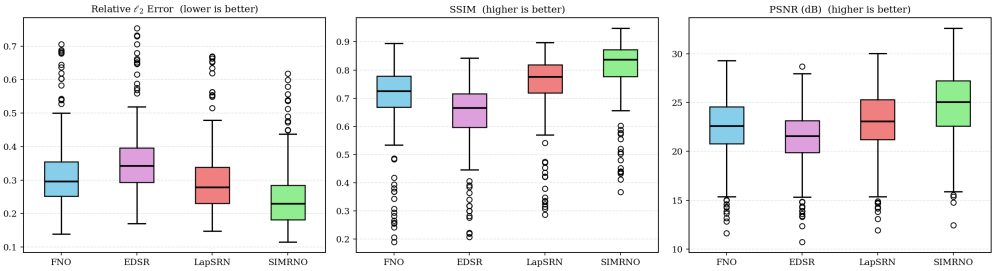
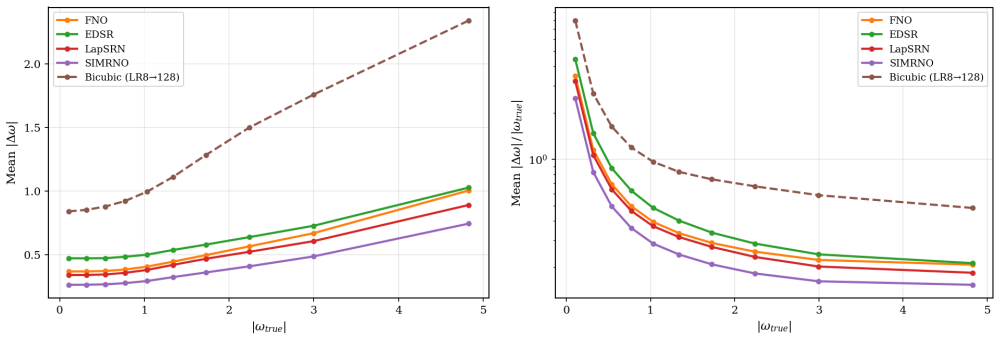


Fig. 2: Kolmogorov-forced turbulence: test-set distributions of relative ℓ_2 error, SSIM, and PSNR for the learned models. Boxes span the interquartile range, the black line marks the median, whiskers extend to 1.5 times the interquartile range, and circles mark individual snapshots beyond the whiskers.

6.1.2 Error conditioned on local vorticity

Figure 3 resolves the reconstruction error by the local magnitude of the reference vorticity (Eq. (44)). The absolute error of every method grows with $|\omega|$, so that the largest errors occur in vortex cores. The ranking of the methods is the same in every bin: SIMR-NO, LapSRN, FNO, EDSR, and bicubic interpolation, in order of increasing error. The mean absolute error of SIMR-NO grows from about 0.26 in the quiescent background to about 0.75 in the strongest vortices, compared with about 0.89 (LapSRN), 1.01 (FNO), 1.03 (EDSR), and 2.34 (bicubic) in the highest bin, so that its advantage in absolute terms is largest where the vorticity is most intense. Relative to the bin-averaged vorticity magnitude (Figure 3b), the error of all methods exceeds unity in the weakest-vorticity bin, where any absolute error is large in relative terms, and decreases to about 0.15 (SIMR-NO), 0.18 (LapSRN), 0.21 (FNO and EDSR), and 0.48 (bicubic) in vortex cores. SIMR-NO therefore reconstructs both the weak background and the strong vortices more accurately than the baselines.



(a) Mean absolute error

(b) Mean absolute error relative to $|\omega|$

Fig. 3: Kolmogorov-forced turbulence: reconstruction error conditioned on the local reference vorticity magnitude, computed over all test pixels grouped into ten quantile bins of $|\omega|$ (Eq. (44)). Left to right along each axis: from the quiescent background to the vortex cores. The vertical axis of (b) is logarithmic.

6.1.3 Representative reconstructions

Figures 4–6 show three test snapshots with their spectral diagnostics, and Table 3 lists the corresponding per-sample metrics. The snapshots were selected among those with a clear advantage of SIMR-NO, to illustrate where its reconstructions differ from those of the baselines; Table 2 gives the unbiased test-set summary. In each field figure, the first row shows the 8×8 input, the pseudo-high-resolution field, the reference, and the bicubic reconstruction with its absolute error; the following rows show the reconstructions and absolute errors of FNO, EDSR, LapSRN, and SIMR-NO. All absolute-error panels of a figure share one colour scale, set by the 99.5th percentile of the errors of the four learned models, and larger errors saturate.

In all three cases, the bicubic field reproduces only the largest structures of the reference (74–79% relative error) and misses the vortex rings and filaments. FNO recovers the large-scale organization but superimposes oriented, grid-scale striations, visible as granular patterns in its error maps. EDSR and LapSRN produce smooth fields whose errors concentrate in the cores of the strongest vortices, in particular at the vortex that crosses the periodic boundary near the top of the domain in Cases 1 and 2. SIMR-NO reproduces the vortex rings, the filaments connecting them, and the core amplitudes most closely, and its errors are the lowest of all methods.

In Case 1 (Figure 4), SIMR-NO reduces the relative ℓ_2 error to 16.56%, compared with 26.40% for FNO, the most accurate baseline in this snapshot, a reduction of 37.3%; its MSE is 60.7% lower and its PSNR 4.05 dB higher than those of FNO. In Case 2 (Figure 5), the baselines show their largest errors in the core of the central vortex ring and at the boundary-crossing vortex; SIMR-NO attains 13.47% against 21.03% for LapSRN (35.9% lower), with a 59.0% lower MSE and a PSNR gain of 3.87 dB. In Case 3 (Figure 6), which contains two vortex rings of opposite sign, LapSRN misrepresents both ring cores; SIMR-NO attains 13.34% against 24.12% for FNO (44.7% lower), with a 69.4% lower MSE and a PSNR gain of 5.14 dB.

The spectral diagnostics in Figures 4–6(b) show the same behaviour in wavenumber space. Bicubic interpolation overestimates the energy at $k = 2$–3 and underestimates both spectra at intermediate wavenumbers, $6 \lesssim k \lesssim 15$. All learned models follow the reference to within a factor of about two up to $k \approx 12$. Beyond $k \approx 15$, FNO develops a spurious high-wavenumber plateau that lies about four decades above the reference at the highest wavenumbers shown and corresponds to the striations in its fields, and EDSR also retains excess high-wavenumber content. The high-wavenumber spectra of SIMR-NO lie closest to the steep decay of the reference in all three cases, although every learned model, including SIMR-NO, retains more small-scale content than the reference beyond $k \approx 20$.

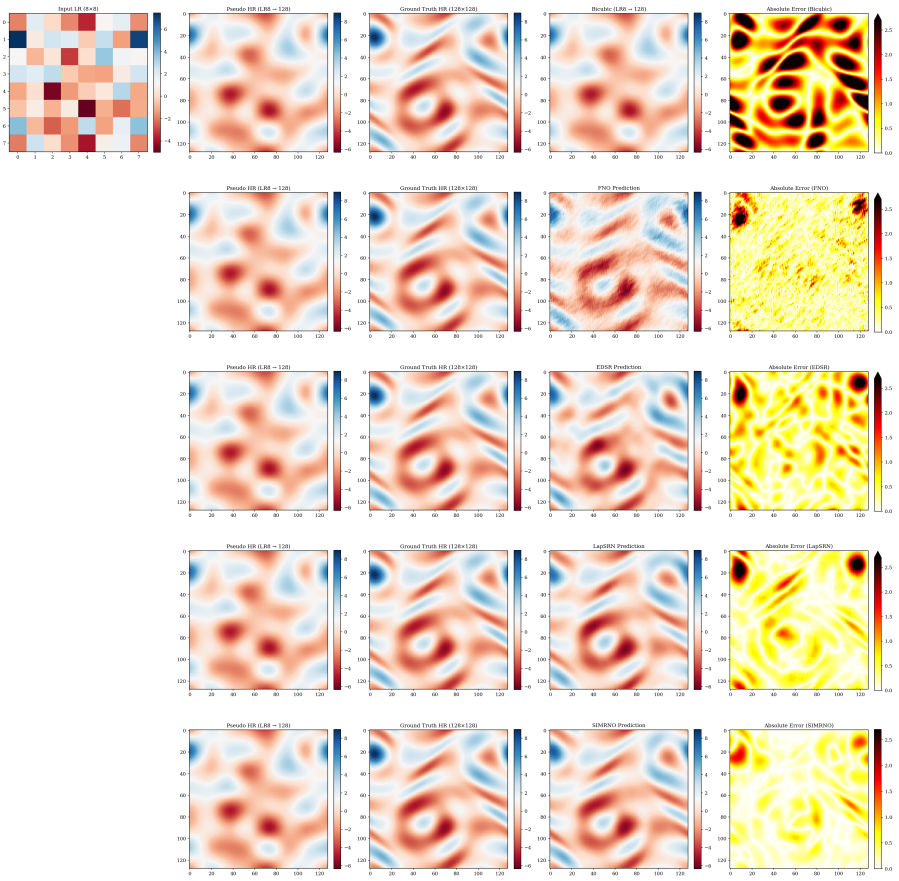
6.1.4 Localized comparison

Figure 7 magnifies a 20×20 region in the lower-right corner of a test snapshot. The box is covered by at most four samples of the 8×8 input, so the vortical structure inside it must be inferred almost entirely by the reconstruction. Within this region (Eq. (45)), SIMR-NO attains a relative ℓ_2 error of 6.46%, compared with 21.18% (FNO), 22.12% (EDSR), 24.08% (LapSRN), and 77.20% (bicubic), that is, about one third of the error of the best baseline. LapSRN, the most accurate baseline over the full test set, is the least accurate learned model in this region, which shows that full-field averages can conceal large local errors near individual structures. These values are regional errors and are not directly comparable with the full-field errors of Table 2.

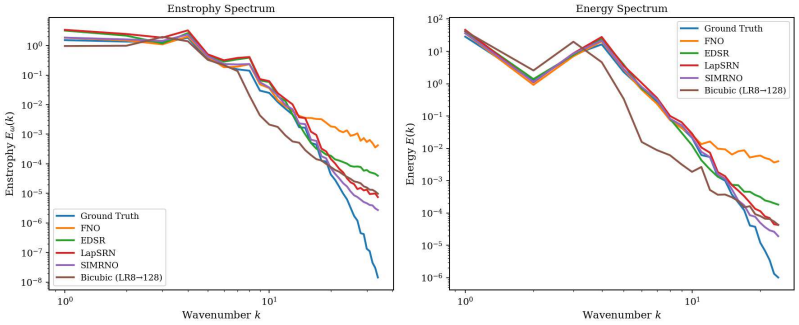
6.2 Decaying Homogeneous Isotropic Turbulence

6.2.1 Aggregate performance and Reynolds-number dependence

Table 4 lists the test-set metrics for DHIT at $Re = 1000$, 4000, and 8000, and Figure 8 shows their dependence on the Reynolds number. At every Reynolds number, bicubic interpolation yields a mean relative ℓ_2 error of about 125% and an SSIM below 0.01,

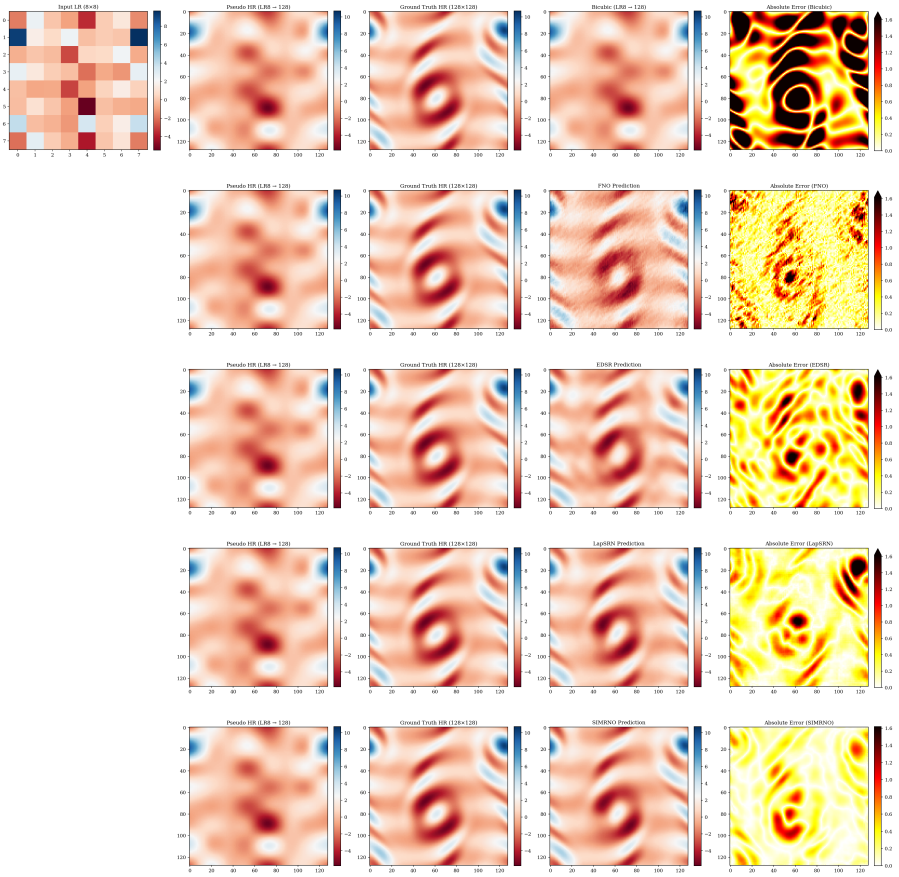


(a) Vorticity fields and absolute errors

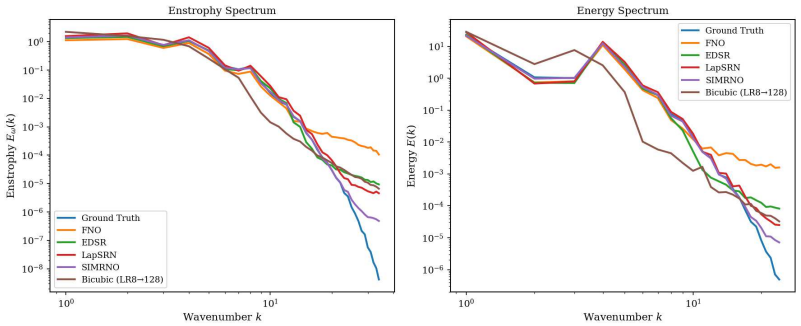


(b) Normalised spectral diagnostics (Section 5)

Fig. 4: Kolmogorov-forced turbulence, Case 1. (a) First row: 8×8 input, pseudo-high-resolution field, reference, bicubic reconstruction, and its absolute error; following rows: pseudo-high-resolution field, reference, reconstruction, and absolute error for FNO, EDSR, LapSRN, and SIMR-NO. (b) Normalised enstrophy and energy diagnostics of the reference and all reconstructions.

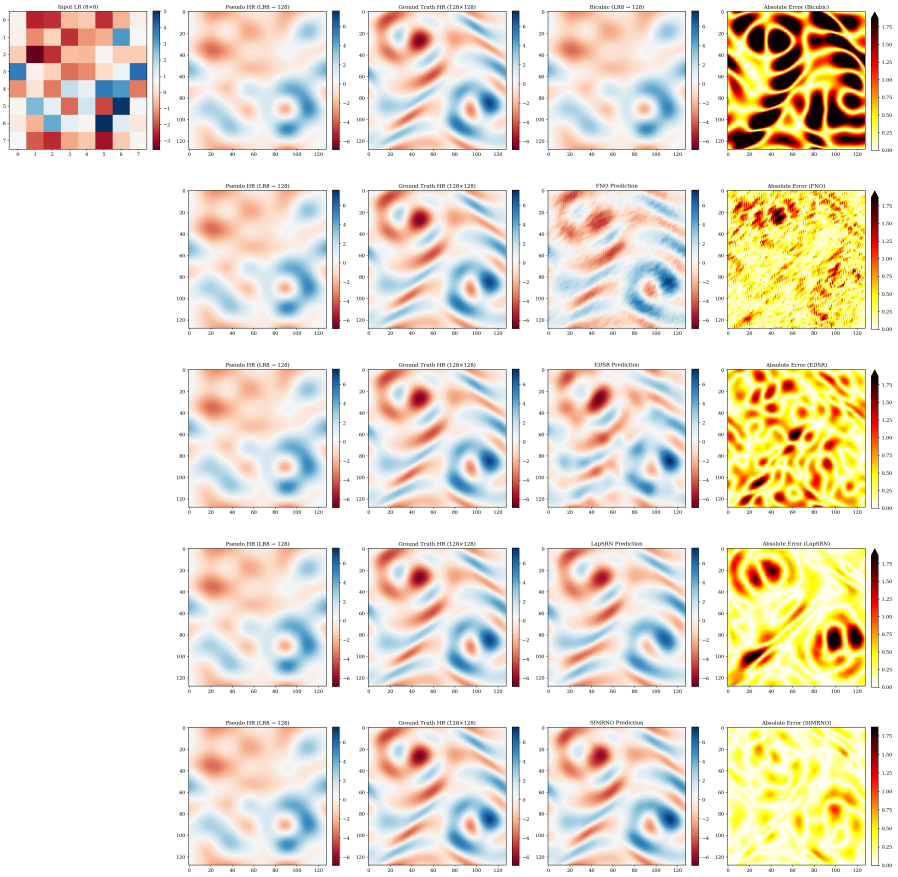


(a) Vorticity fields and absolute errors

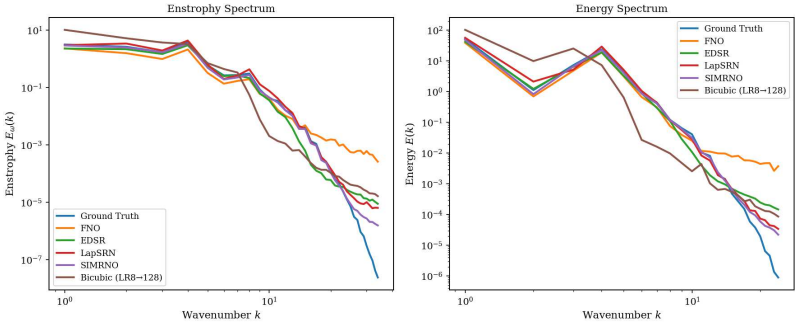


(b) Normalised spectral diagnostics

Fig. 5: Kolmogorov-forced turbulence, Case 2, with the layout of Figure 4. The baselines show their largest errors in the core of the central vortex ring and at the vortex crossing the periodic boundary.



(a) Vorticity fields and absolute errors



(b) Normalised spectral diagnostics

Fig. 6: Kolmogorov-forced turbulence, Case 3, with the layout of Figure 4. The snapshot contains two vortex rings of opposite sign, whose cores dominate the errors of the baselines.

Table 3: Kolmogorov-forced turbulence: per-sample metrics of the three representative test snapshots shown in Figures 4–6. The best value in each row is shown in bold.

Case	Metric	Bicubic	FNO	EDSR	LapSRN	SIMR-NO
1	MSE ↓	2.5587	0.3228	0.4413	0.3525	**0.1269**
	SSIM ↑	0.2605	0.7857	0.7178	0.7956	**0.8960**
	PSNR (dB) ↑	14.95	23.95	22.59	23.56	**28.00**
	Rel. ℓ_2 (%) ↓	74.34	26.40	30.87	27.59	**16.56**
2	MSE ↓	2.6622	0.2129	0.2387	0.1889	**0.0775**
	SSIM ↑	0.2337	0.8103	0.7716	0.8323	**0.9333**
	PSNR (dB) ↑	15.99	26.96	26.46	27.48	**31.35**
	Rel. ℓ_2 (%) ↓	78.95	22.33	23.64	21.03	**13.47**
3	MSE ↓	2.6572	0.2454	0.3406	0.3273	**0.0751**
	SSIM ↑	0.2068	0.7479	0.7032	0.7765	**0.8905**
	PSNR (dB) ↑	12.66	23.01	21.58	21.76	**28.15**
	Rel. ℓ_2 (%) ↓	79.36	24.12	28.41	27.85	**13.34**

that is, an error larger than that of the trivial prediction $\widehat{\omega} = 0$. The energy-containing vortices of DHIT lie largely at wavenumbers above the Nyquist wavenumber $k_c = 4$ of the coarse grid, so the 8×8 samples are dominated by aliased small-scale content (Section 3.4), and the interpolated field contains large-scale structures with incorrect amplitude and phase. This failure is insensitive to Re (124.9–125.5%), because it is set mainly by the position of the energy-containing scales relative to k_c.

The learned models separate into two groups at every Reynolds number. The convolutional networks EDSR and LapSRN reach mean relative errors between 18% and 33%, whereas the operator-learning models FNO and SIMR-NO remain below 8%. SIMR-NO is the most accurate method for every metric and every Reynolds number:

- At $Re = 1000$, SIMR-NO attains 0.0277 ± 0.0101, which is 7.7% lower than FNO (0.0300) and 86.5% and 85.0% lower than EDSR (0.2049) and LapSRN (0.1850), with a PSNR of 45.23 dB, 0.60 dB above FNO. At this Reynolds number, FNO and SIMR-NO are close in all three metrics.
- At $Re = 4000$, SIMR-NO attains 0.0316 ± 0.0134, which is 45.6% lower than FNO (0.0581) and 89.3% and 89.2% lower than EDSR (0.2946) and LapSRN (0.2917), with a PSNR of 45.36 dB, 4.85 dB above FNO.
- At $Re = 8000$, SIMR-NO attains 0.0318 ± 0.0190, which is 54.8% lower than FNO (0.0704) and 90.2% and 90.3% lower than EDSR (0.3248) and LapSRN (0.3274), with a PSNR of 46.16 dB, 7.00 dB above FNO.

The baselines become less accurate as the Reynolds number increases, whereas SIMR-NO is nearly unaffected. From $Re = 1000$ to $Re = 8000$, the mean relative error increases by a factor of 2.35 for FNO ($0.0300 \to 0.0704$), 1.59 for EDSR, and 1.77 for LapSRN, but only by a factor of 1.15 for SIMR-NO ($0.0277 \to 0.0318$), with almost no change between $Re = 4000$ and $Re = 8000$. Correspondingly, the PSNR

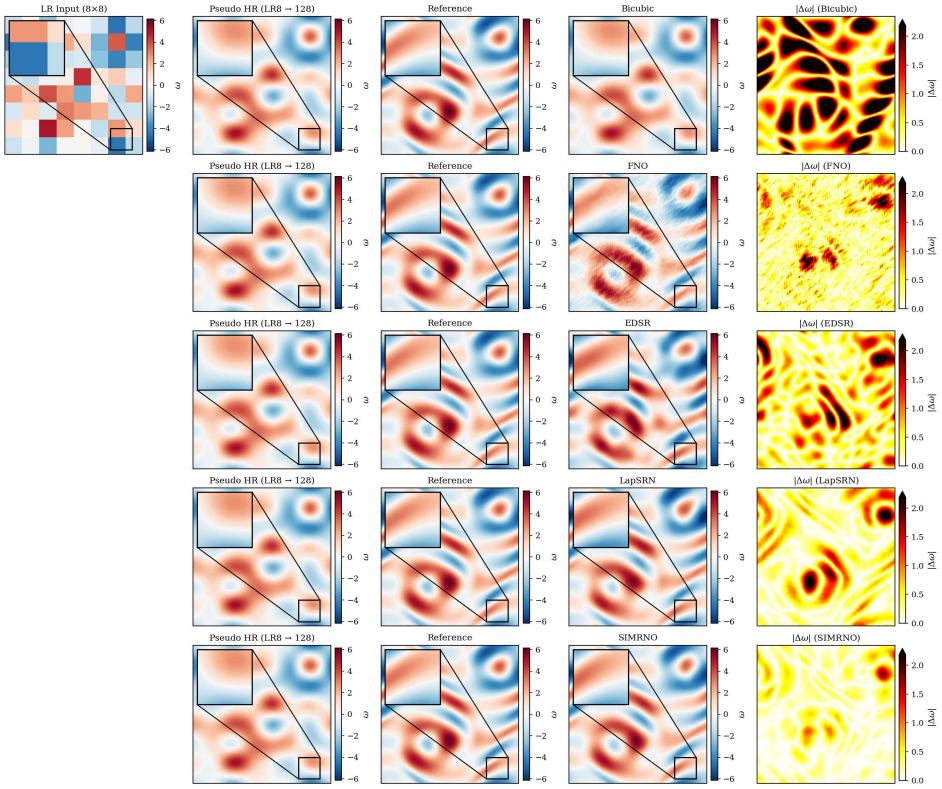


Fig. 7: Kolmogorov-forced turbulence: localized comparison. Each field panel shows the full 128×128 field with a 20×20 box (lower right) and its magnified view (upper left); the first column shows the 8×8 input with the same box. The last column shows the full-field absolute error. Relative ℓ_2 error inside the box: bicubic 77.20%, FNO 21.18%, EDSR 22.12%, LapSRN 24.08%, SIMR-NO **6.46**%.

of FNO decreases by 5.46 dB, that of EDSR by 2.38 dB, and that of LapSRN by 3.33 dB, whereas the PSNR of SIMR-NO remains between 45.23 and 46.16 dB. The SSIM of EDSR and LapSRN decreases from about 0.95 to 0.89 and 0.87, while that of SIMR-NO stays at or above 0.998. As a consequence, the advantage of SIMR-NO over FNO, the strongest baseline at every Reynolds number, grows with Re: the error ratio FNO/SIMR-NO is 1.08, 1.84, and 2.21 at $Re = 1000$, 4000, and 8000, respectively. At $Re = 1000$, where viscosity keeps the field smooth, the two operator-learning models are therefore nearly equivalent. The standard deviation of the relative error of SIMR-NO is smaller than that of FNO at every Reynolds number, by a factor of 1.3, 2.8, and 2.1, so that its accuracy also varies less across the decay. Among the convolutional networks, LapSRN is more accurate than EDSR at $Re = 1000$, whereas the two are nearly indistinguishable at $Re = 4000$ and 8000, where EDSR attains the higher SSIM.

This trend is consistent with the physics of the decaying flow. A larger Reynolds number reduces the viscous damping of Eq. (6), so that the enstrophy cascade extends

Table 4: Decaying homogeneous isotropic turbulence: test-set metrics (mean ± standard deviation) at $Re = 1000$, 4000, and 8000. Arrows indicate whether lower or higher values are better; the best value for each Reynolds number is shown in bold.

Re	Method	Rel. ℓ_2 ↓	SSIM ↑	PSNR (dB) ↑
1000	Bicubic	1.2523 ± 0.0451	0.0076 ± 0.0070	11.59 ± 0.82
	FNO	0.0300 ± 0.0129	0.9982 ± 0.0021	44.63 ± 3.13
	EDSR	0.2049 ± 0.0276	0.9474 ± 0.0205	27.41 ± 1.20
	LapSRN	0.1850 ± 0.0242	0.9475 ± 0.0190	28.29 ± 1.19
	SIMR-NO	$\mathbf{0.0277 \pm 0.0101}$	$\mathbf{0.9985 \pm 0.0016}$	$\mathbf{45.23 \pm 2.76}$
4000	Bicubic	1.2554 ± 0.0635	0.0055 ± 0.0064	12.68 ± 1.48
	FNO	0.0581 ± 0.0374	0.9929 ± 0.0126	40.51 ± 3.95
	EDSR	0.2946 ± 0.0460	0.9042 ± 0.0326	25.39 ± 1.28
	LapSRN	0.2917 ± 0.0435	0.8884 ± 0.0368	25.46 ± 1.23
	SIMR-NO	$\mathbf{0.0316 \pm 0.0134}$	$\mathbf{0.9984 \pm 0.0019}$	$\mathbf{45.36 \pm 2.85}$
8000	Bicubic	1.2489 ± 0.0617	0.0049 ± 0.0048	13.18 ± 1.81
	FNO	0.0704 ± 0.0391	0.9897 ± 0.0111	39.16 ± 3.88
	EDSR	0.3248 ± 0.0545	0.8855 ± 0.0402	25.03 ± 1.34
	LapSRN	0.3274 ± 0.0533	0.8671 ± 0.0441	24.95 ± 1.31
	SIMR-NO	$\mathbf{0.0318 \pm 0.0190}$	$\mathbf{0.9983 \pm 0.0027}$	$\mathbf{46.16 \pm 3.52}$

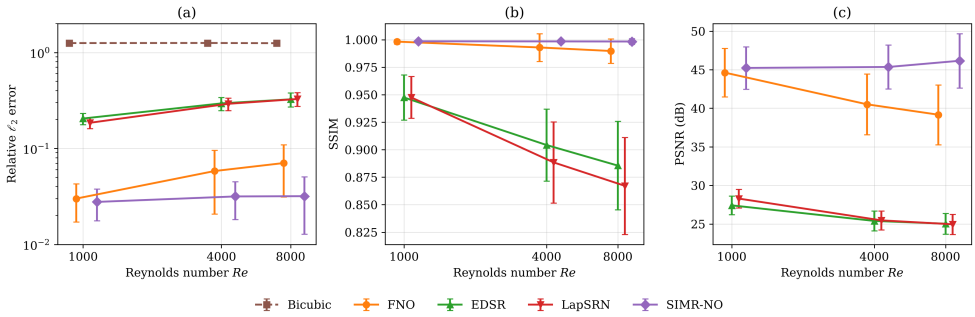


Fig. 8: Decaying homogeneous isotropic turbulence: test-set (a) relative ℓ_2 error (logarithmic axis), (b) SSIM, and (c) PSNR versus Reynolds number (values of Table 4). Markers show the mean and error bars one standard deviation over the 81 test snapshots; markers are offset horizontally for legibility. Bicubic interpolation is shown only in (a) because its SSIM (< 0.01) and PSNR (11.6–13.2 dB) lie far outside the range of the learned models.

to higher wavenumbers, vorticity gradients steepen, and a larger fraction of the vorticity variance resides near the truncation wavenumber $k = 64$ of the target grid (Section 3.1.2). At $Re = 1000$, viscosity smooths the field and the fine-scale content that must be inferred is comparatively small; at $Re = 8000$, the reconstruction must

recover thin filaments and sharp vortex edges that are entirely unresolved by the 8×8 observations. SIMR-NO retains this content with almost no loss of accuracy, whereas FNO and the convolutional networks degrade substantially; the present experiments do not isolate which architectural component is responsible for this difference.

These errors are substantially lower than for Kolmogorov flow. Part of this difference reflects the data partition (Section 3.3): at each Reynolds number, the DHIT snapshots are drawn from a single trajectory at intervals of 0.01, so test snapshots are temporally close to training snapshots. The comparison between methods is unaffected, because all models are trained and tested on identical partitions.

6.2.2 Representative reconstructions at $Re = 1000$

Figures 9–11 show three test snapshots at $Re = 1000$ from the late ($t = 4.00$), intermediate ($t = 2.94$), and early ($t = 0.28$) stages of the decay, and Table 5 lists their metrics. These are the same snapshot indices as Cases 1–3 at $Re = 4000$, so the two Reynolds numbers can be compared at fixed times. As the flow decays, the field evolves from a dense population of small vortices at $t = 0.28$ to fewer, larger, coherent vortices connected by filaments at $t = 4.00$; at this Reynolds number, viscosity keeps these filaments comparatively broad.

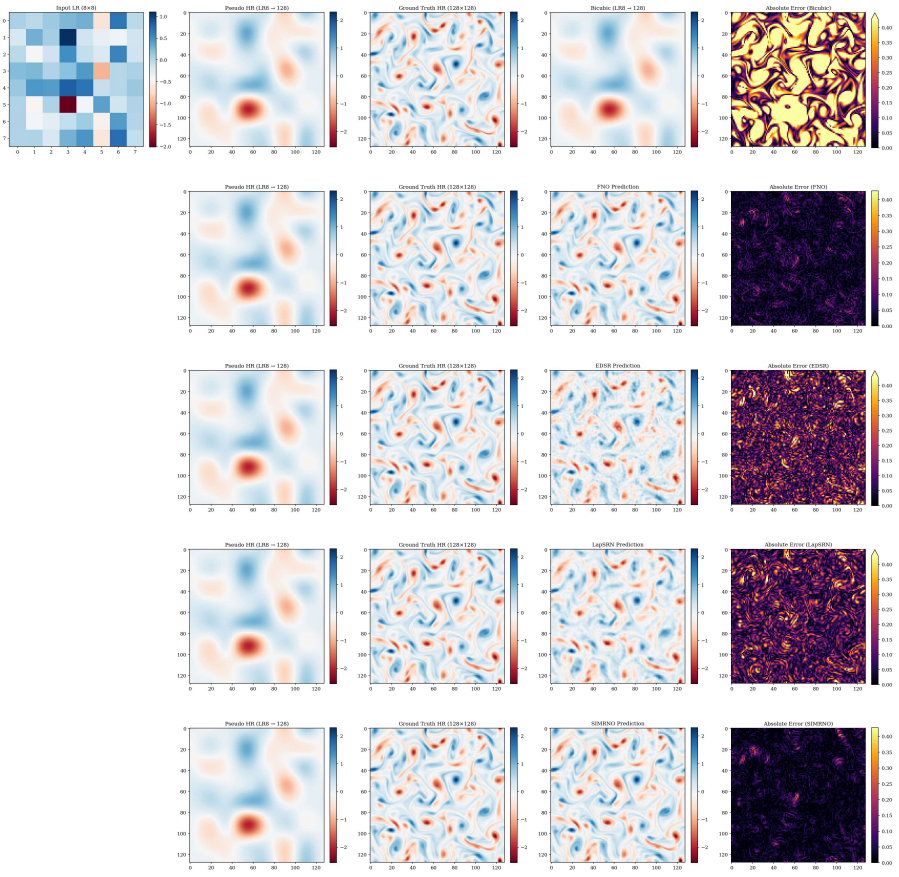
Bicubic interpolation reproduces no vortices and exceeds 120% relative error in all three snapshots. EDSR and LapSRN recover the positions of the vortices, with relative errors between 18% and 31% that appear as bright speckles along vortex edges and filaments. FNO reaches 5.0–10.2%, and SIMR-NO has the lowest error in every case: 8.30% against 10.23% for FNO at $t = 4.00$ (18.9% lower, +1.82 dB), 4.54% against 5.53% at $t = 2.94$ (18.0% lower, +1.72 dB), and 2.91% against 4.96% at $t = 0.28$ (41.4% lower, +4.64 dB). The margin over FNO is therefore largest in the early-time snapshot, whose fine-scale vortices are the most numerous, and smaller at late times, when the field is dominated by a few smooth, large vortices.

The spectra show the same ranking. Bicubic interpolation exceeds the reference at $k \leq 3$, by up to several decades, and falls below it at all larger wavenumbers. The spectral peak moves from $k \approx 13$ at $t = 0.28$ to $k \approx 4$ at $t = 4.00$. At high wavenumbers, the enstrophy spectra of EDSR and LapSRN fall below the reference beyond $k \approx 30$–35, by up to about 0.7 decade at $k = 64$ in the early and intermediate snapshots, whereas FNO and SIMR-NO remain close to the reference. At the largest scales, $k = 1$–2, which carry a small fraction of the energy, LapSRN overestimates the enstrophy of the early-time snapshot by more than one decade and EDSR and FNO by a smaller margin; SIMR-NO lies closest to the reference at these wavenumbers.

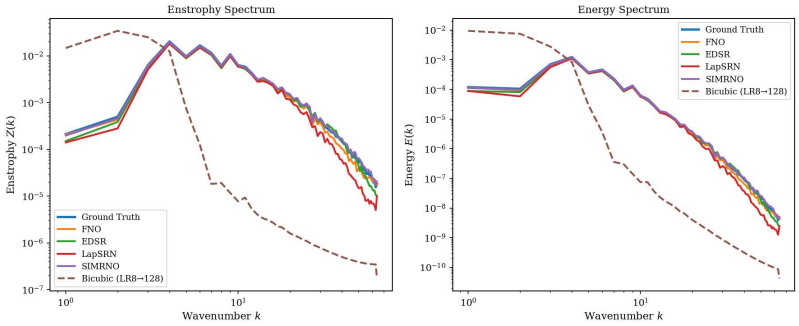
6.2.3 Representative reconstructions at $Re = 4000$

Figures 12–14 show the same three test snapshots as at $Re = 1000$, from the late ($t = 4.00$), intermediate ($t = 2.94$), and early ($t = 0.28$) stages of the decay, and Table 6 lists their metrics. Compared with $Re = 1000$, the vortices are more compact and the filaments connecting them are thinner at every time.

In all three snapshots, bicubic interpolation produces a smooth large-scale field with no discernible vortices, and its error saturates the colour scale. EDSR and LapSRN recover the positions of the vortices but with speckled errors throughout the

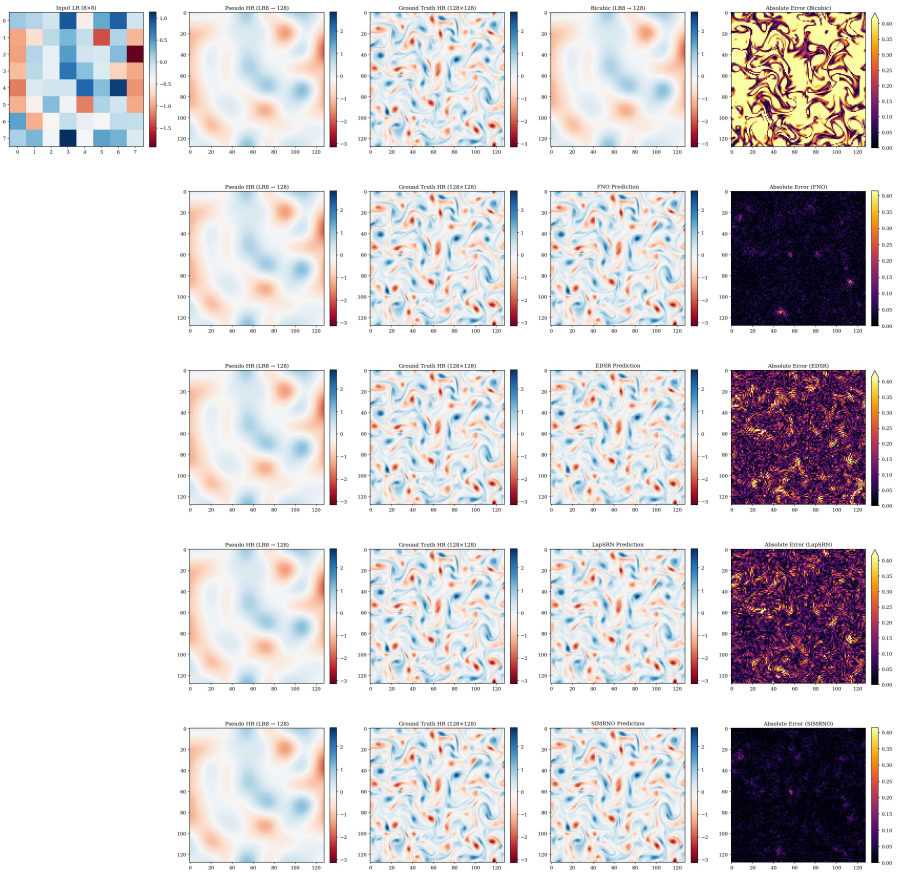


(a) Vorticity fields and absolute errors

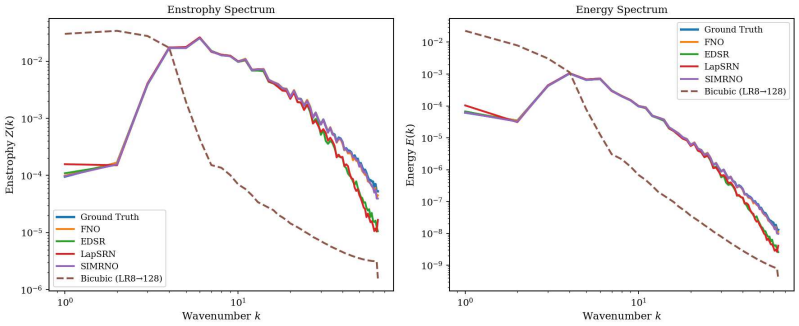


(b) Enstrophy spectrum $\mathcal{E}(k)$ and energy spectrum $E(k)$

Fig. 9: Decaying homogeneous isotropic turbulence at $Re = 1000$, Case 1 (late stage, $t = 4.00$), with the layout of Figure 4. (b) Spectra are computed with Eqs. (14)–(15) over $1 \leq k \leq 64$ without rescaling.

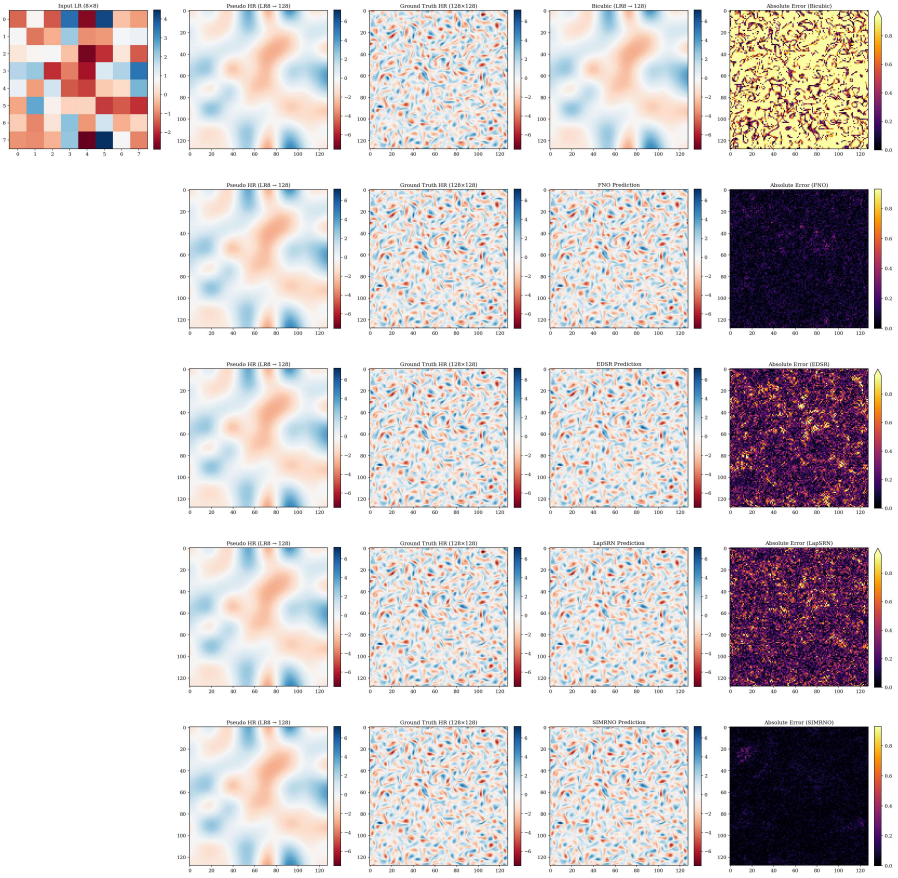


(a) Vorticity fields and absolute errors

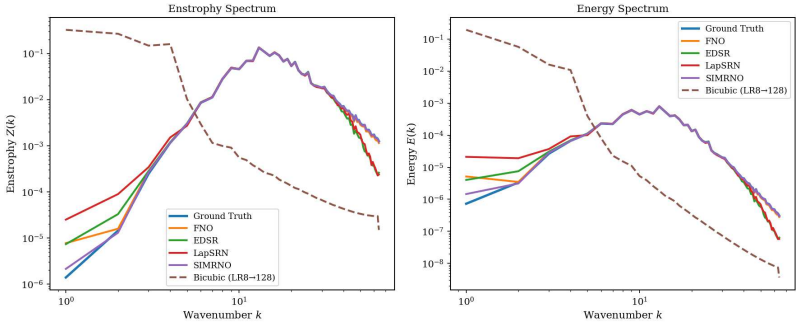


(b) Enstrophy spectrum $\mathcal{E}(k)$ and energy spectrum $E(k)$

Fig. 10: Decaying homogeneous isotropic turbulence at $Re = 1000$, Case 2 (intermediate stage, $t = 2.94$), with the layout of Figure 9. EDSR and LapSRN underestimate the enstrophy at the highest wavenumbers, whereas SIMR-NO follows the reference.



(a) Vorticity fields and absolute errors



(b) Enstrophy spectrum $\mathcal{E}(k)$ and energy spectrum $E(k)$

Fig. 11: Decaying homogeneous isotropic turbulence at $Re = 1000$, Case 3 (early stage, $t = 0.28$), with the layout of Figure 9. The baselines overestimate the largest-scale enstrophy, whereas SIMR-NO follows the reference over the full wavenumber range.

Table 5: Decaying homogeneous isotropic turbulence at $Re = 1000$: per-sample metrics of the three test snapshots shown in Figures 9–11. The best value in each row is shown in bold.

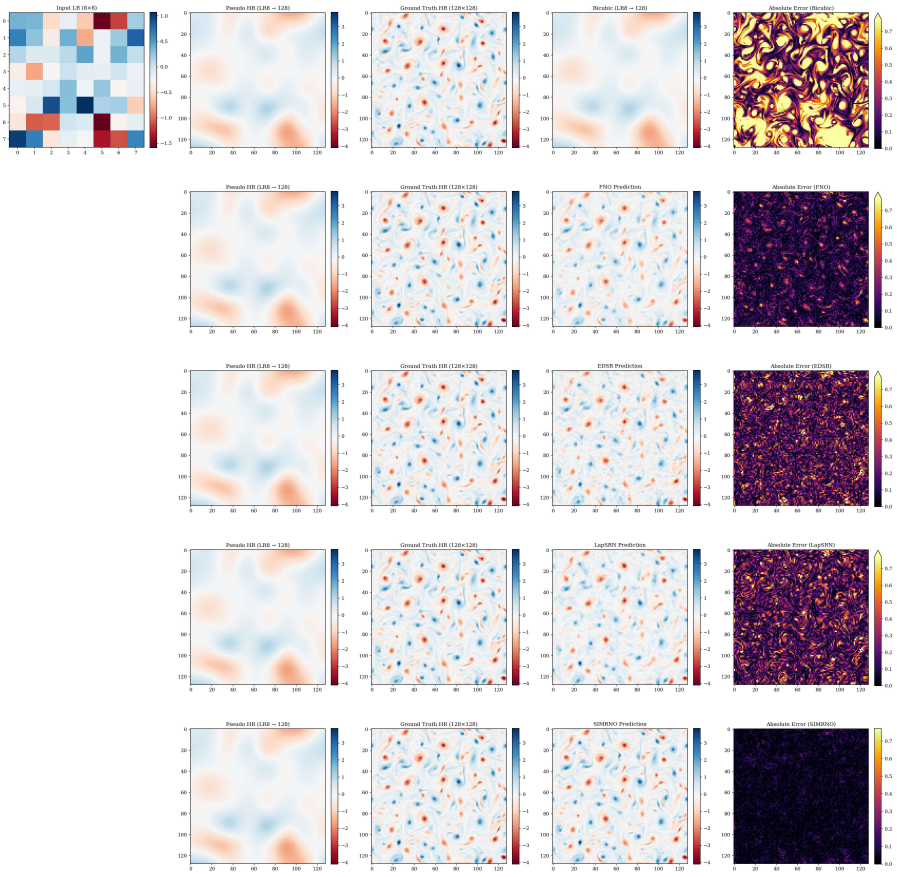
Case	Metric	Bicubic	FNO	EDSR	LapSRN	SIMR-NO
1 $t = 4.00$	MSE ↓	0.3913	0.0027	0.0245	0.0196	**0.0018**
	SSIM ↑	0.0372	0.9817	0.8302	0.8432	**0.9865**
	PSNR (dB) ↑	12.17	33.81	24.20	25.18	**35.62**
	Rel. ℓ_2 (%) ↓	123.54	10.23	30.91	27.62	**8.30**
2 $t = 2.94$	MSE ↓	0.5814	0.0012	0.0227	0.0190	**0.0008**
	SSIM ↑	0.0105	0.9951	0.9151	0.9171	**0.9965**
	PSNR (dB) ↑	12.13	39.00	26.22	26.99	**40.72**
	Rel. ℓ_2 (%) ↓	121.89	5.53	24.07	22.02	**4.54**
3 $t = 0.28$	MSE ↓	4.8732	0.0075	0.1118	0.0988	**0.0026**
	SSIM ↑	0.0016	0.9961	0.9619	0.9539	**0.9988**
	PSNR (dB) ↑	10.46	38.60	26.85	27.39	**43.24**
	Rel. ℓ_2 (%) ↓	126.67	4.96	19.18	18.03	**2.91**

domain, and FNO concentrates its errors in individual vortex cores. The error maps of SIMR-NO are nearly uniform and close to zero. In the late-time snapshot (Case 1, Figure 12), SIMR-NO attains a relative ℓ_2 error of 9.20% against 24.79% for FNO, a reduction of 62.9%, with a PSNR gain of 8.61 dB; EDSR and LapSRN exceed 41%. In the intermediate snapshot (Case 2, Figure 13), the error of SIMR-NO is 7.41% against 19.21% for FNO (61.4% lower, +8.28 dB). In the early-time snapshot (Case 3, Figure 14), whose fine-scale vortices are the most numerous, SIMR-NO attains 3.41% against 10.40% for FNO (67.2% lower, +9.69 dB). At the same three times, the margin of SIMR-NO over FNO is thus considerably larger than at $Re = 1000$ (18–41%).

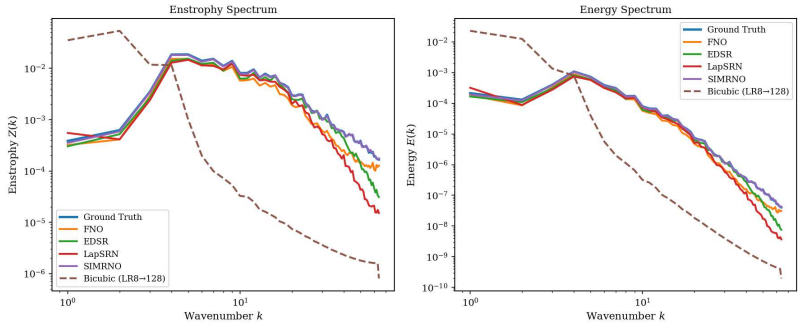
The spectra in Figures 12–14(b) show the spectral peak moving from $k \approx 13$ at $t = 0.28$ to $k \approx 4$ at $t = 4.00$. The spectrum of the bicubic field exceeds the reference at $k \leq 3$, where the aliased coarse samples deposit their variance, and falls one to three decades below it at all smaller scales. At high wavenumbers, the enstrophy spectra of EDSR and LapSRN fall below the reference beyond $k \approx 25$–35, by up to about one decade at $k = 64$, consistent with the loss of fine-scale content in their fields; FNO underestimates the highest wavenumbers by a smaller margin, most visibly at late times. At the largest scales, $k = 1$–2, the baselines can deviate strongly: in the early-time snapshot, LapSRN and FNO overestimate the enstrophy at $k = 1$ by more than one decade. SIMR-NO follows the reference spectra across the full wavenumber range in all three snapshots.

6.2.4 Representative reconstructions at $Re = 8000$

Figures 15–17 show three test snapshots at $Re = 8000$, from the early ($t = 0.28$), intermediate ($t = 1.22$), and late ($t = 4.00$) stages of the decay, and Table 7 lists their metrics. Cases 1 and 3 are the same snapshot indices as Cases 3 and 1 at $Re = 1000$

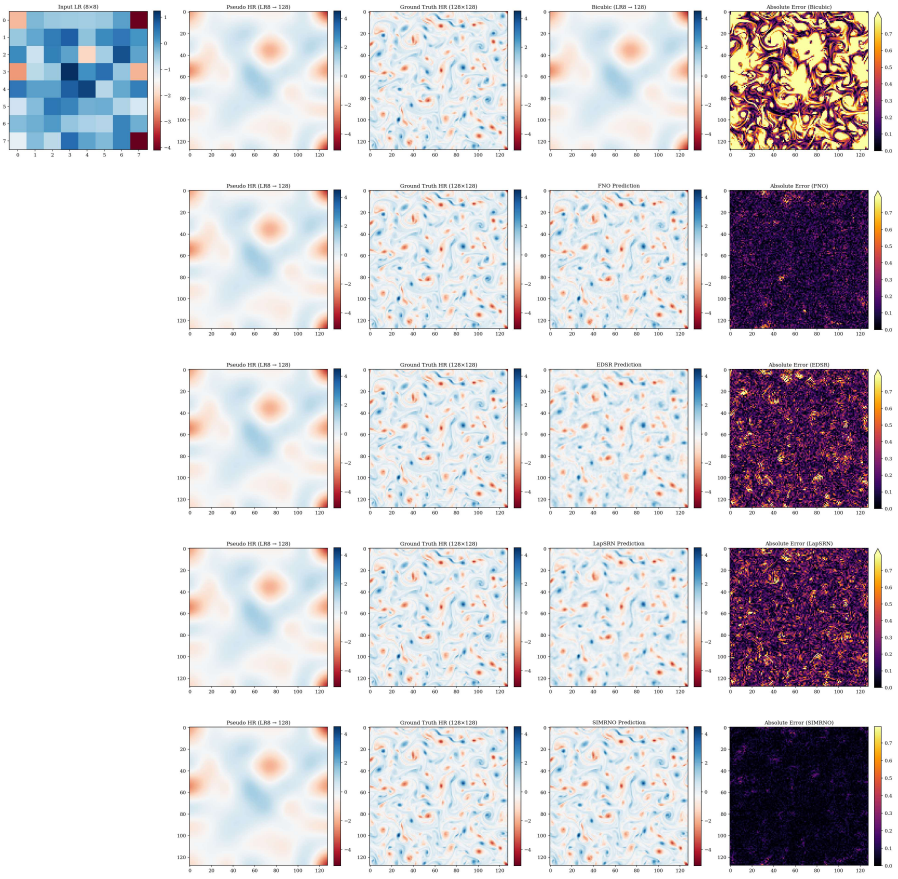


(a) Vorticity fields and absolute errors

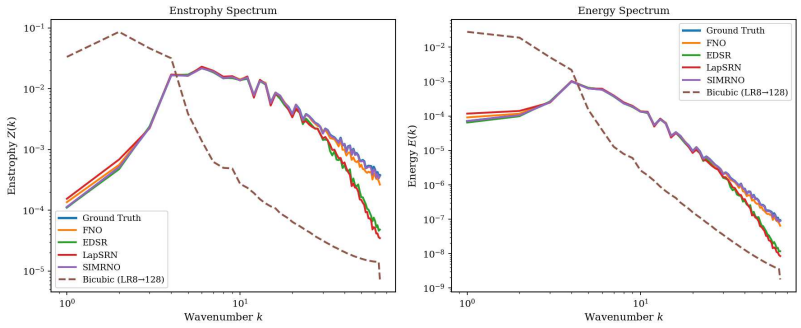


(b) Enstrophy spectrum $\mathcal{E}(k)$ and energy spectrum $E(k)$

Fig. 12: Decaying homogeneous isotropic turbulence at $Re = 4000$, Case 1 (late stage, $t = 4.00$), with the layout of Figure 9.

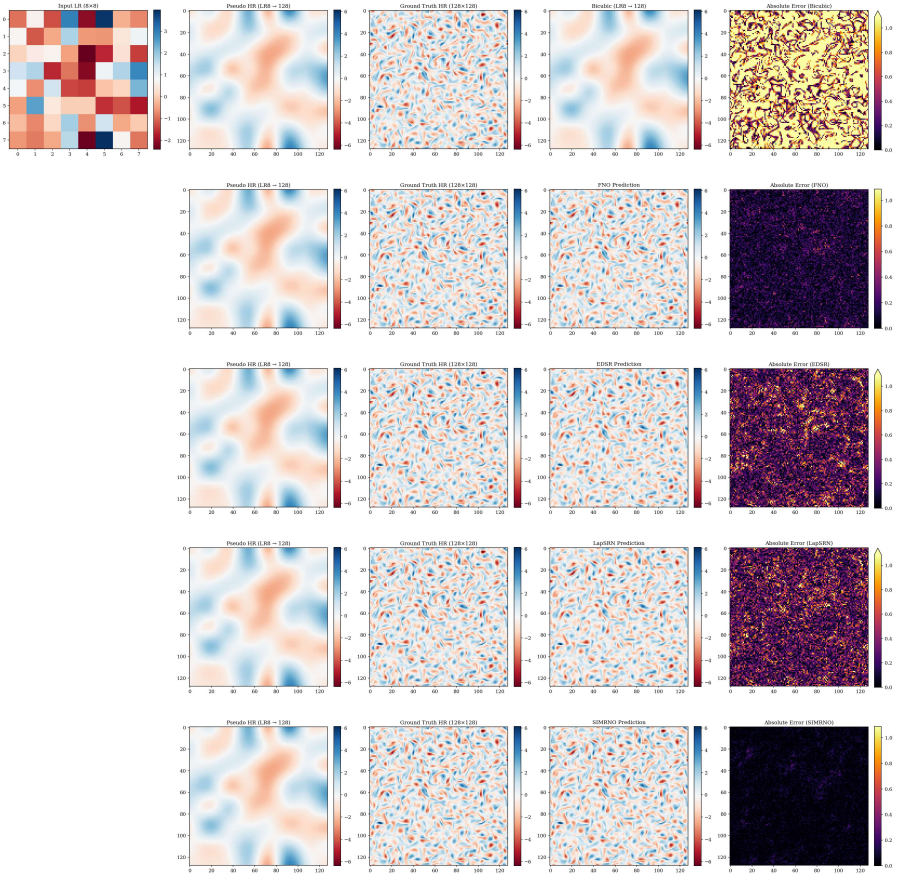


(a) Vorticity fields and absolute errors

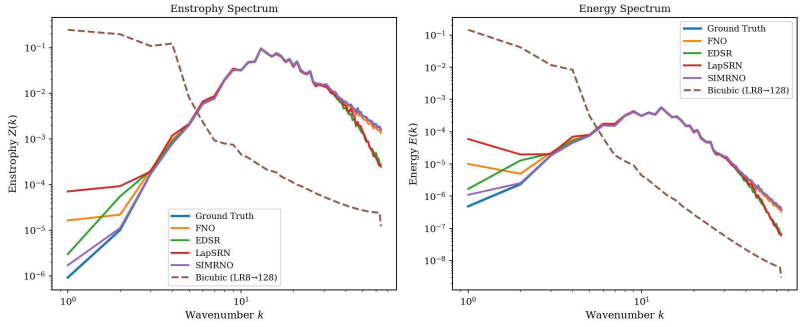


(b) Enstrophy spectrum $\mathcal{E}(k)$ and energy spectrum $E(k)$

Fig. 13: Decaying homogeneous isotropic turbulence at $Re = 4000$, Case 2 (intermediate stage, $t = 2.94$), with the layout of Figure 9.



(a) Vorticity fields and absolute errors



(b) Enstrophy spectrum $\mathcal{E}(k)$ and energy spectrum $E(k)$

Fig. 14: Decaying homogeneous isotropic turbulence at $Re = 4000$, Case 3 (early stage, $t = 0.28$), with the layout of Figure 9. The baselines overestimate the largest-scale enstrophy, whereas SIMR-NO follows the reference over the full wavenumber range.

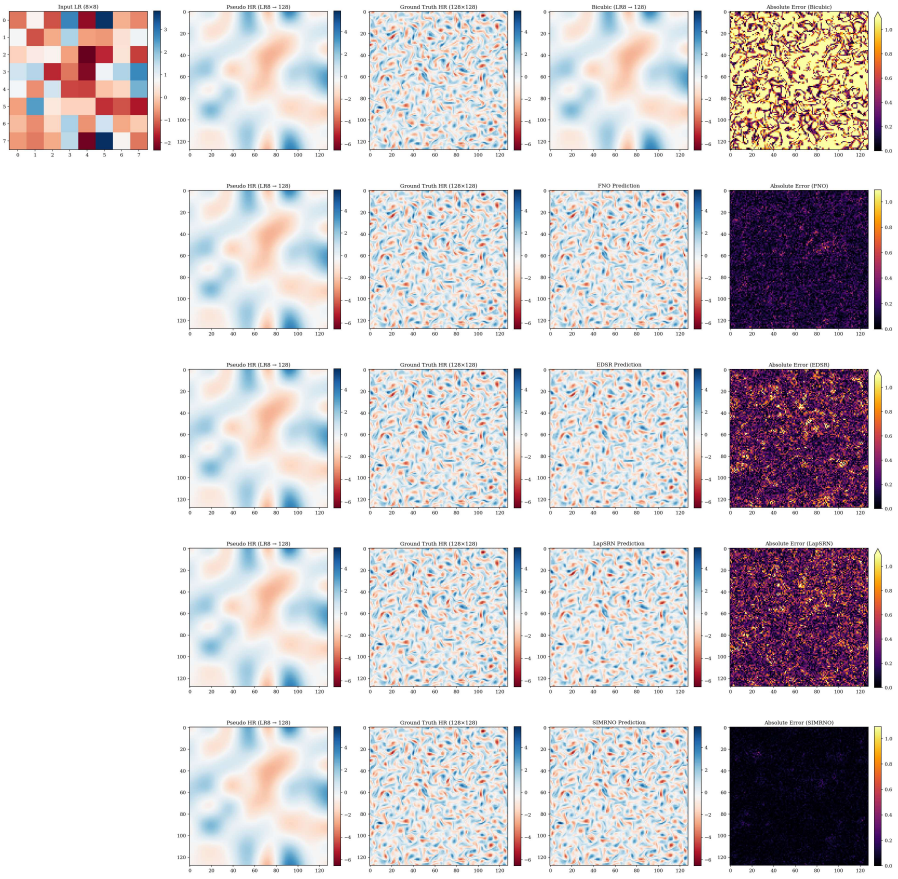
Table 6: Decaying homogeneous isotropic turbulence at $Re = 4000$: per-sample metrics of the three test snapshots shown in Figures 12–14. The best value in each row is shown in bold.

Case	Metric	Bicubic	FNO	EDSR	LapSRN	SIMR-NO
1 $t = 4.00$	MSE ↓	0.6228	0.0250	0.0691	0.0720	**0.0034**
	SSIM ↑	0.0028	0.9009	0.7480	0.6783	**0.9849**
	PSNR (dB) ↑	14.30	28.26	23.85	23.68	**36.87**
	Rel. ℓ_2 (%) ↓	123.62	24.79	41.19	42.02	**9.20**
2 $t = 2.94$	MSE ↓	0.9516	0.0202	0.0718	0.0689	**0.0030**
	SSIM ↑	0.0128	0.9541	0.8405	0.8329	**0.9937**
	PSNR (dB) ↑	13.30	30.02	24.52	24.70	**38.30**
	Rel. ℓ_2 (%) ↓	131.75	19.21	36.20	35.45	**7.41**
3 $t = 0.28$	MSE ↓	3.6758	0.0250	0.1511	0.1489	**0.0027**
	SSIM ↑	0.0011	0.9811	0.9331	0.9124	**0.9986**
	PSNR (dB) ↑	10.40	32.09	24.27	24.33	**41.78**
	Rel. ℓ_2 (%) ↓	126.27	10.40	25.60	25.41	**3.41**

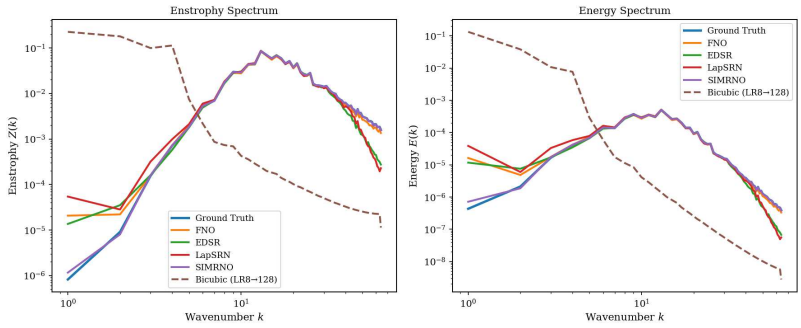
and $Re = 4000$, which allows a direct comparison across Reynolds numbers at fixed times.

In the early-time snapshot (Case 1), SIMR-NO attains 3.20% against 11.06% for FNO, a reduction of 71.1% with a PSNR gain of 10.78 dB, while EDSR and LapSRN exceed 26%. At $t = 0.28$, the error of FNO increases from 4.96% at $Re = 1000$ to 10.40% at $Re = 4000$ and 11.06% at $Re = 8000$, whereas that of SIMR-NO remains between 2.91% and 3.41%. In the intermediate snapshot (Case 2), SIMR-NO attains 4.89% against 12.53% for FNO (61.0% lower, +8.18 dB), and EDSR and LapSRN exceed 35%. The late-time snapshot (Case 3) is the most difficult case for every method: the field consists of intense, isolated coherent vortices connected by thin filaments, and EDSR and LapSRN reach 44.17% and 43.12% with an SSIM below 0.74. SIMR-NO attains 10.17% against 13.38% for FNO, a smaller but consistent reduction of 24.0% (+2.39 dB). At every Reynolds number, the errors of all learned models are highest at $t = 4.00$, and at $Re = 8000$ the margin of SIMR-NO over FNO is the smallest of the three cases.

In the spectra (Figures 15–17(b)), the deficit of the convolutional networks at high wavenumbers is pronounced in all three snapshots: the enstrophy spectra of EDSR and LapSRN fall below the reference beyond $k \approx 25$–35, by up to about one decade at $k = 64$, and FNO underestimates the highest wavenumbers by a smaller margin. At $k = 1$–2 in the early-time snapshot, LapSRN, FNO, and EDSR overestimate the reference, whose energy at these wavenumbers is several orders of magnitude below the spectral peak. SIMR-NO follows the reference spectra up to $k = 64$ in all three snapshots.

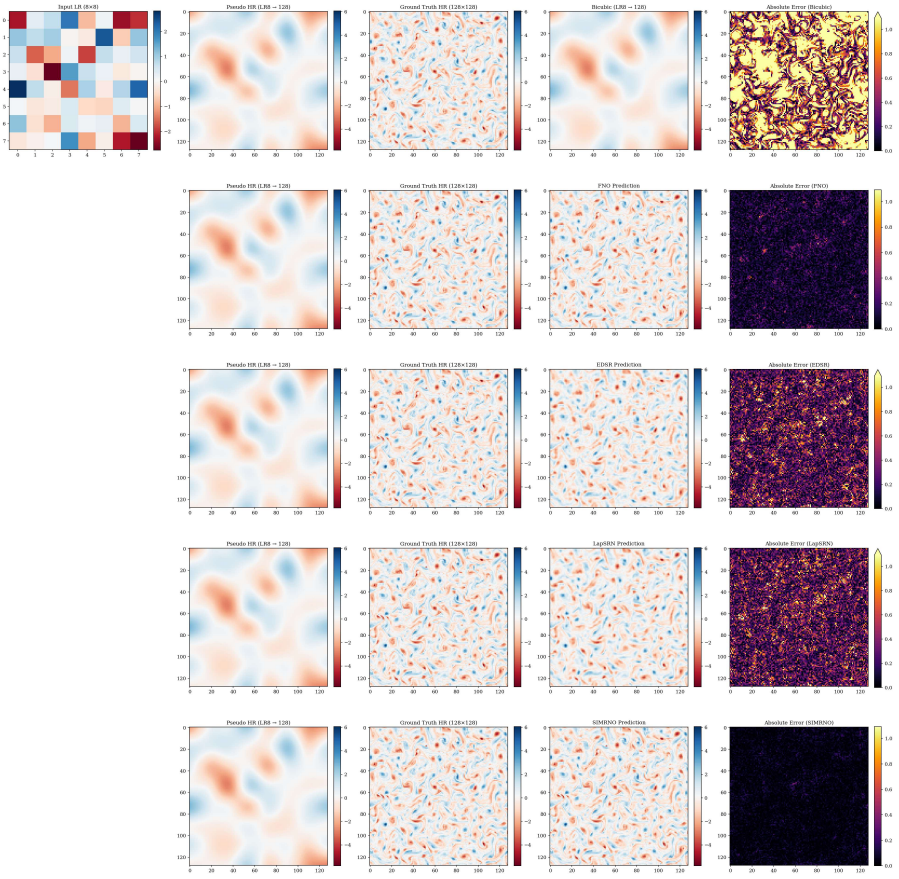


(a) Vorticity fields and absolute errors

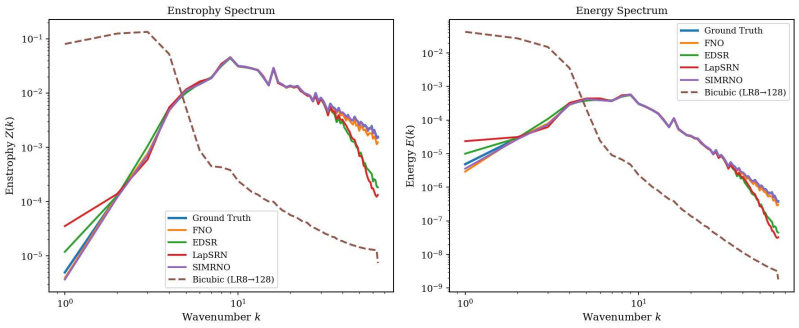


(b) Enstrophy spectrum $\mathcal{E}(k)$ and energy spectrum $E(k)$

Fig. 15: Decaying homogeneous isotropic turbulence at $Re = 8000$, Case 1 (early stage, $t = 0.28$), with the layout of Figure 9.

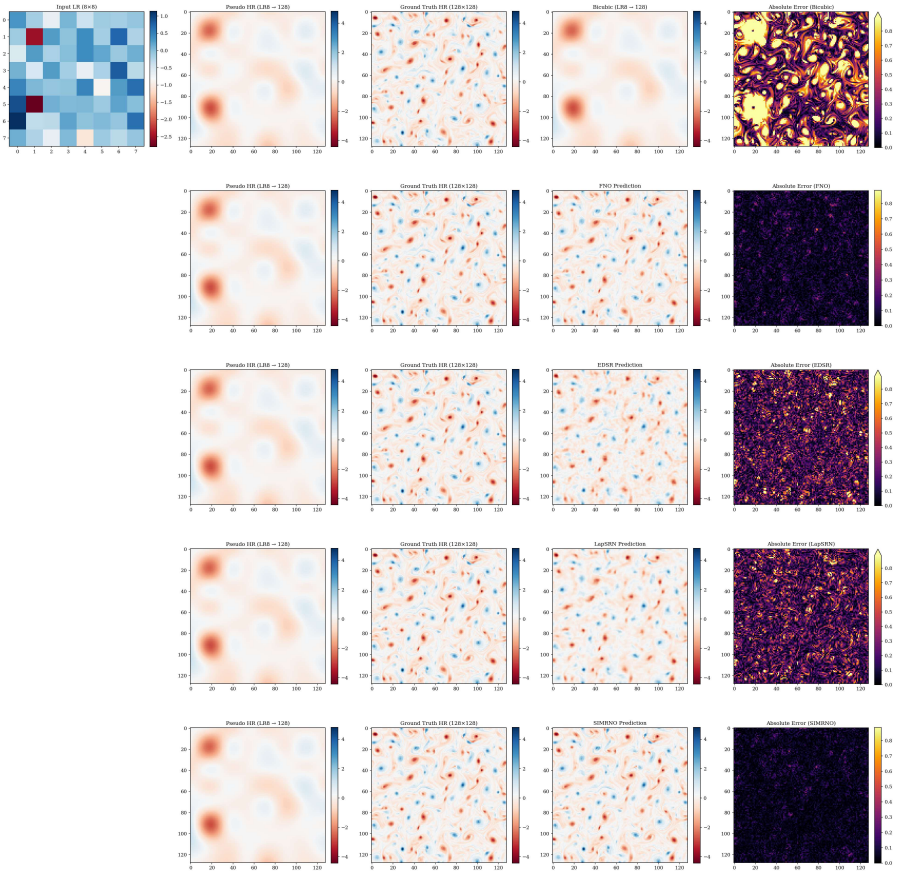


(a) Vorticity fields and absolute errors

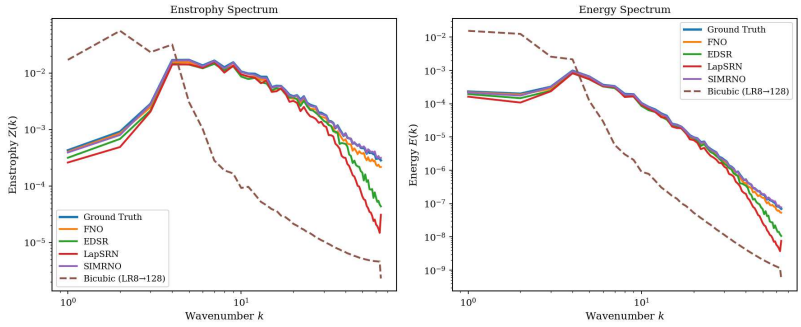


(b) Enstrophy spectrum $\mathcal{E}(k)$ and energy spectrum $E(k)$

Fig. 16: Decaying homogeneous isotropic turbulence at $Re = 8000$, Case 2 (intermediate stage, $t = 1.22$), with the layout of Figure 9.



(a) Vorticity fields and absolute errors



(b) Enstrophy spectrum $\mathcal{E}(k)$ and energy spectrum $E(k)$

Fig. 17: Decaying homogeneous isotropic turbulence at $Re = 8000$, Case 3 (late stage, $t = 4.00$), with the layout of Figure 9. EDSR and LapSRN fall up to about one decade below the reference enstrophy at the highest wavenumbers, whereas SIMR-NO follows the reference.

Table 7: Decaying homogeneous isotropic turbulence at $Re = 8000$: per-sample metrics of the three test snapshots shown in Figures 15–17. The best value in each row is shown in bold.

Case	Metric	Bicubic	FNO	EDSR	LapSRN	SIMR-NO
1 $t = 0.28$	MSE ↓	3.3777	0.0259	0.1530	0.1587	**0.0022**
	SSIM ↑	0.0010	0.9737	0.9230	0.9077	**0.9986**
	PSNR (dB) ↑	10.38	31.53	23.82	23.66	**42.31**
	Rel. ℓ_2 (%) ↓	126.21	11.06	26.86	27.35	**3.20**
2 $t = 1.22$	MSE ↓	1.9490	0.0185	0.1508	0.1488	**0.0028**
	SSIM ↑	0.0039	0.9823	0.8793	0.8684	**0.9975**
	PSNR (dB) ↑	12.75	32.96	23.86	23.92	**41.14**
	Rel. ℓ_2 (%) ↓	128.49	12.53	35.74	35.50	**4.89**
3 $t = 4.00$	MSE ↓	0.7011	0.0081	0.0884	0.0843	**0.0047**
	SSIM ↑	0.0153	0.9734	0.7374	0.7018	**0.9828**
	PSNR (dB) ↑	15.19	34.56	24.19	24.39	**36.94**
	Rel. ℓ_2 (%) ↓	124.35	13.38	44.17	43.12	**10.17**

6.2.5 Localized comparisons

Figures 18–20 magnify a 20×20 region in the lower-right corner of one test snapshot at each Reynolds number. Each region is covered by at most four samples of the 8×8 input, so its vortical structure must be inferred almost entirely by the reconstruction. Table 8 lists the local relative ℓ_2 errors of Eq. (45).

At $Re = 1000$ and $t = 4.00$ (Figure 18), the region contains broad, viscously smoothed vorticity bands. SIMR-NO attains a local error of 8.83%, compared with 13.38% for FNO and about 31–35% for LapSRN and EDSR. At $Re = 4000$ and $t = 0.83$ (Figure 19), the region contains small vortices of both signs connected by curved filaments; SIMR-NO attains 4.98%, compared with 15.18% for FNO and about 32% for EDSR and LapSRN. At $Re = 8000$ and $t = 0.83$ (Figure 20), the corresponding region contains thinner filaments; SIMR-NO attains 6.04%, whereas the error of FNO is 19.07% and those of EDSR and LapSRN are 37–40%. The local error ratio FNO/SIMR-NO thus increases from 1.5 at $Re = 1000$ to 3.0 at $Re = 4000$ and 3.2 at $Re = 8000$, in line with the full-field trend of Figure 8. These are regional errors and are not directly comparable with the full-field errors of Table 4.

6.3 Comparison of the Two Turbulence Regimes

The two flows lead to the same conclusion about the best method but to different rankings of the baselines. In Kolmogorov flow, the three baselines lie within a narrow range (30–36%), LapSRN is the strongest, and SIMR-NO improves on it by 17.7% in mean relative error. In DHIT, FNO is the strongest baseline at every Reynolds number and SIMR-NO improves on it by 7.7%, 45.6%, and 54.8% at $Re = 1000$, 4000, and 8000, while the two convolutional networks remain between 18% and 33%. The convolutional baselines therefore perform comparatively well when much of the

Table 8: Decaying homogeneous isotropic turbulence: relative ℓ_2 error (%) inside the 20×20 box of the localized comparisons (Figures 18–20). The best value in each row is shown in bold.

Re	t	Bicubic	FNO	EDSR	LapSRN	SIMR-NO
1000	4.00	126.57	13.38	34.82	31.33	**8.83**
4000	0.83	135.98	15.18	31.75	32.07	**4.98**
8000	0.83	114.58	19.07	36.57	39.77	**6.04**

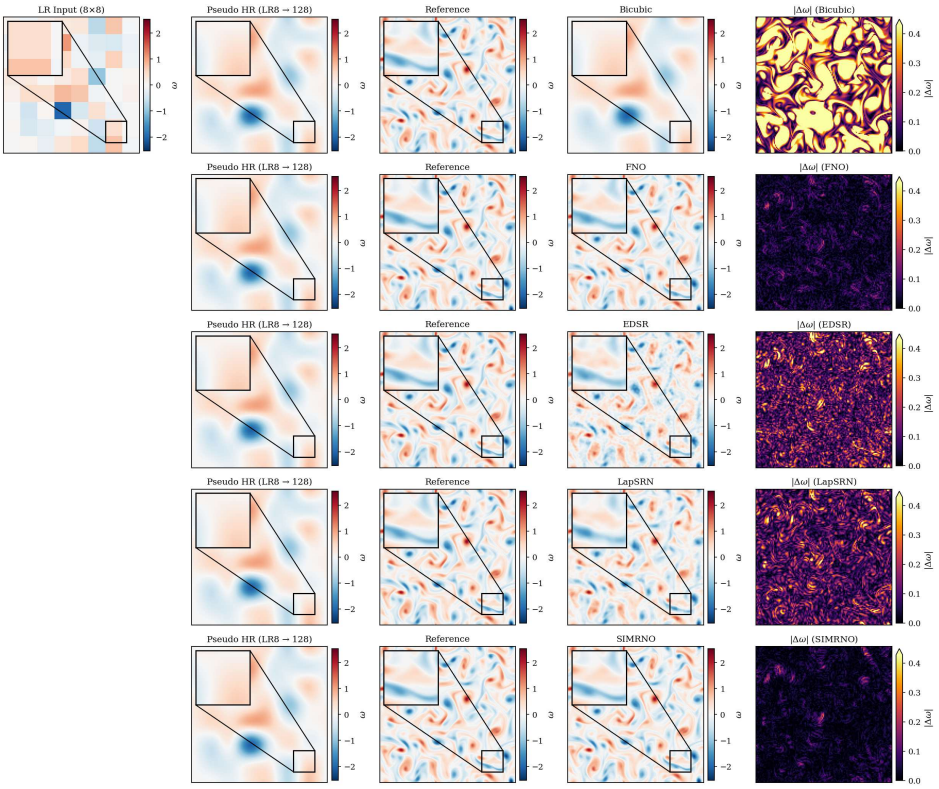


Fig. 18: Decaying homogeneous isotropic turbulence at $Re = 1000$: localized comparison for the test snapshot at $t = 4.00$, with the layout of Figure 7. Relative ℓ_2 error inside the box: bicubic 126.57%, FNO 13.38%, EDSR 34.82%, LapSRN 31.33%, SIMR-NO **8.83**%.

vorticity variance resides at the forcing scale $k_f = 4$, which the coarse grid partially resolves, and comparatively poorly when the vortical content lies above the coarse-grid Nyquist wavenumber, as in DHIT, whereas the operator-learning models gain most in

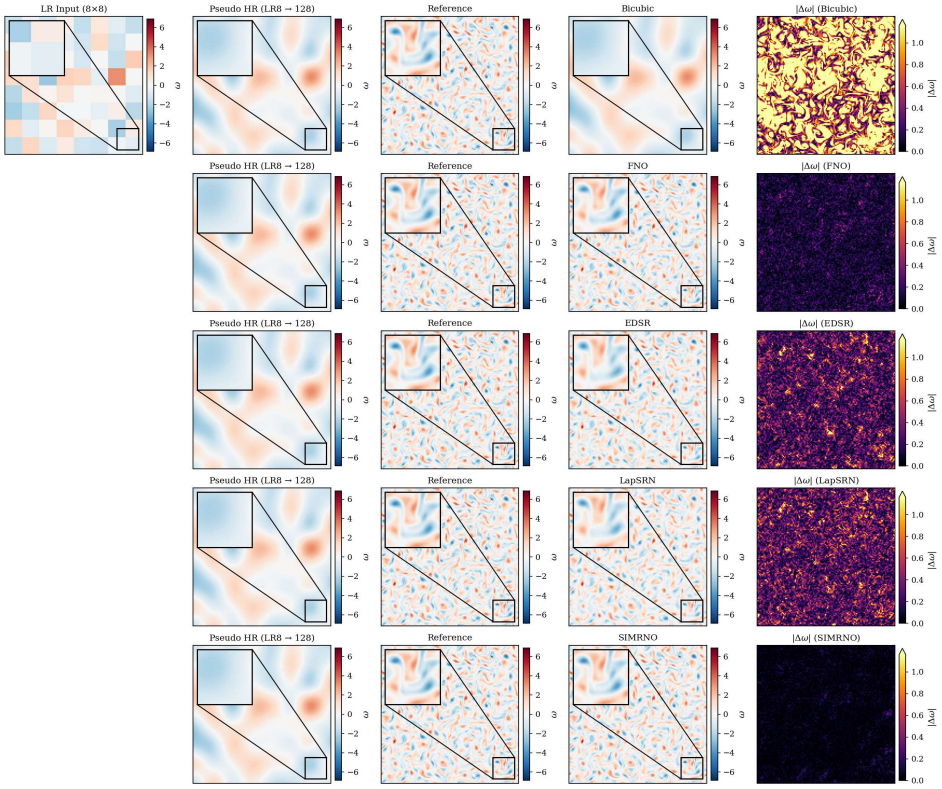


Fig. 19: Decaying homogeneous isotropic turbulence at $Re = 4000$: localized comparison for the test snapshot at $t = 0.83$, with the layout of Figure 7. Relative ℓ_2 error inside the box: bicubic 135.98%, FNO 15.18%, EDSR 31.75%, LapSRN 32.07%, SIMR-NO **4.98**%.

the latter case. Only SIMR-NO is the most accurate method in both regimes, which is consistent with its combination of global spectral corrections and local refinement.

In both regimes, the advantage of SIMR-NO is largest where reconstruction is most difficult: in the cores of strong vortices for Kolmogorov flow (Figure 3); in localized structures that are covered by a single coarse sample (Figures 7 and 18–20), where its regional error is about one third of that of the best baseline in Kolmogorov flow and in DHIT at $Re = 4000$ and 8000; and, for DHIT, at the higher Reynolds numbers, where the fine-scale content that must be inferred is largest (Figure 8). Its reconstructions also follow the reference spectral decay more closely than those of the baselines, without the spurious high-wavenumber content of FNO in Kolmogorov flow or the high-wavenumber deficit of EDSR and LapSRN in DHIT.

6.4 Summary of Results

On Kolmogorov-forced turbulence, SIMR-NO attains a mean relative ℓ_2 error of 0.2506 ± 0.0983, which is 17.7%, 21.5%, and 30.8% lower than LapSRN, FNO, and

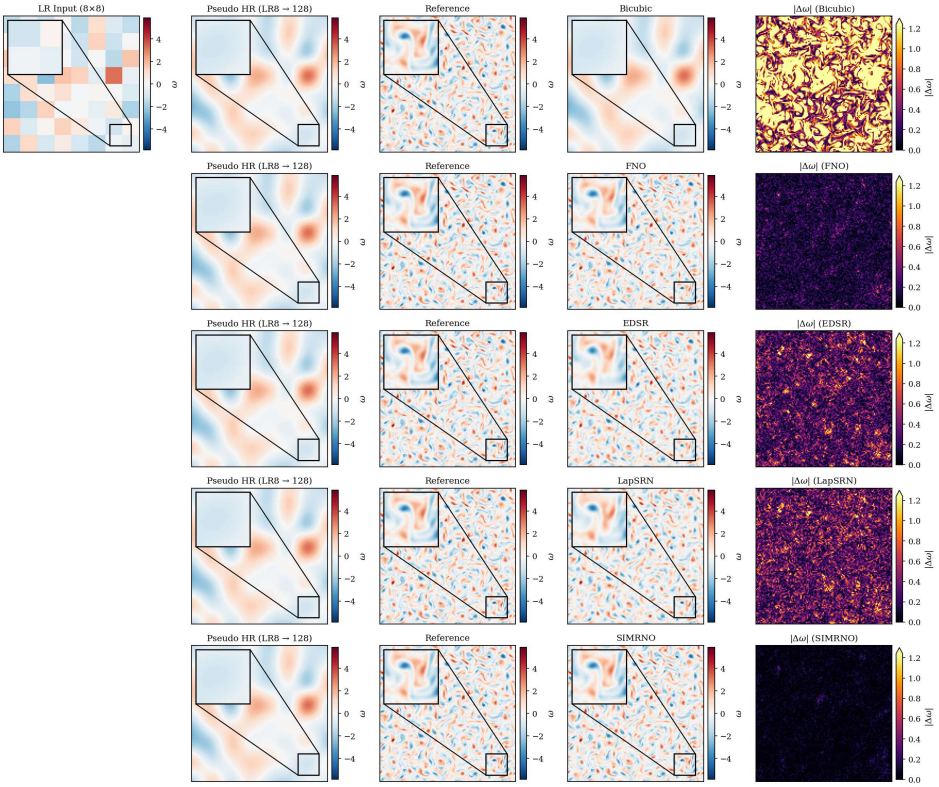


Fig. 20: Decaying homogeneous isotropic turbulence at $Re = 8000$: localized comparison for the test snapshot at $t = 0.83$, with the layout of Figure 7. Relative ℓ_2 error inside the box: bicubic 114.58%, FNO 19.07%, EDSR 36.57%, LapSRN 39.77%, SIMR-NO **6.04**%.

EDSR, and 68% lower than bicubic interpolation, together with the highest SSIM (0.8075) and PSNR (24.77 dB); in a 20×20 region of a test snapshot its error is 6.46%, against 21–24% for the baselines. On decaying homogeneous isotropic turbulence, SIMR-NO attains mean relative errors of 0.0277, 0.0316, and 0.0318 at $Re = 1000$, 4000, and 8000, which are 7.7%, 45.6%, and 54.8% lower than FNO and 85–90% lower than EDSR and LapSRN, with the highest SSIM (≥ 0.998) and PSNR (45.23–46.16 dB) at every Reynolds number. From $Re = 1000$ to 8000, its mean error increases by a factor of 1.15, compared with 1.59–2.35 for the learned baselines. The representative reconstructions, vorticity-conditioned errors, localized comparisons, and spectral diagnostics show that these aggregate gains are accompanied by a more accurate reconstruction of vortex cores, fine-scale filaments, and high-wavenumber content. We hypothesize that these gains result from the combination of the hierarchical multi-resolution factorization, which divides the 16× inverse problem into two 2× residual corrections; the spectrally gated Fourier layers, which weight spectral corrections by wavenumber; and the local refinement modules, which recover fine-scale features

beyond the truncated Fourier basis. Isolating the contribution of each component requires ablation experiments, which are left to future work.

7 Conclusion and Future Work

We introduced the *Spectrally-Informed Multi-Resolution Neural Operator* (SIMR-NO), a hierarchical operator learning framework for reconstructing high-resolution turbulent flow fields from extremely coarse 8×8 observations. The central contributions are threefold: a hierarchical factorization of the ill-posed inverse operator across intermediate spatial resolutions that decomposes the 16× super-resolution problem into tractable per-stage residual learning problems; spectrally gated Fourier convolution layers in which a learned function of the radial wavenumber weights the spectral corrections by scale; and lightweight local refinement modules that recover fine-scale spatial features beyond the reach of the truncated Fourier basis. Together, these components enable SIMR-NO to capture both large-scale coherent structures and small-scale turbulent fluctuations within a unified, end-to-end trainable operator learning framework.

SIMR-NO was evaluated on two two-dimensional turbulent flows with distinct dynamics, statistically stationary Kolmogorov-forced turbulence and freely decaying homogeneous isotropic turbulence at $Re = 1000$, 4000, and 8000, and compared with bicubic interpolation, FNO, EDSR, and LapSRN under an identical training and evaluation protocol. On Kolmogorov flow, SIMR-NO reduces the mean relative ℓ_2 error by 17.7%, 21.5%, and 30.8% relative to LapSRN, FNO, and EDSR and attains the highest SSIM and PSNR. On decaying turbulence, it reduces the error by 7.7%, 45.6%, and 54.8% relative to FNO, the strongest baseline in this regime, at $Re = 1000$, 4000, and 8000, and by 85–90% relative to EDSR and LapSRN; its accuracy is the least sensitive to the Reynolds number among all learned models, and its margin over FNO is small at the lowest Reynolds number. Whereas the ranking of the baselines changes between the two regimes, SIMR-NO is the most accurate method in both. Its advantage is largest in the cores of intense vortices and in localized structures unresolved by the coarse observations, and its reconstructions follow the reference spectral decay more closely than those of the baselines. The DHIT results reported here were obtained within a single trajectory at each Reynolds number; evaluation on independent realisations, generalisation across Reynolds numbers with a single trained model, and ablations of the individual architectural components will further establish the generality of these findings.

The research field currently has multiple promising areas that should be explored. The next logical step involves testing SIMR-NO through three-dimensional turbulence simulations, which use irregular geometric shapes and large-eddy simulation methods and geophysical flow simulation methods. The use of physics-informed training objectives together with Navier-Stokes residual penalties and spectral consistency losses will increase the physical accuracy of the model without needing extra labeled information. The system will achieve time-resolved reconstruction through its architecture, which supports spatio-temporal settings while enabling users to predict and assimilate data. The creation of uncertainty-aware extensions through probabilistic learning or diffusion-based operator learning will produce confidence estimates that are required in

turbulent flow analysis. The hierarchical spectrally-informed operator learning framework, which we present here, enables us to solve various inverse problems that arise in medical imaging and climate downscaling and multi-fidelity simulation acceleration. The research shows that using multiscale priors together with spectral inductive biases through neural operator architectures creates an effective method for achieving super-resolution that maintains physical accuracy in complex dynamical systems.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported in part by the AFOSR Grant FA9550-24-1-0327.

Data availability

Data supporting the findings of this study are available from the corresponding author upon request.

References

[1] Tennekes, H., Lumley, J.L.: A First Course in Turbulence. MIT Press, Cambridge, MA (1972)

[2] Pope, S.B.: Turbulent Flows. Cambridge University Press, Cambridge, UK (2000)

[3] Brunton, S.L., Noack, B.R., Koumoutsakos, P.: Machine learning for fluid mechanics. Annual Review of Fluid Mechanics **52**, 477–508 (2020)

[4] Vinuesa, R., Brunton, S.L.: Enhancing computational fluid dynamics with machine learning. Nature Computational Science **2**, 358–366 (2022)

[5] Kramer, B., Willcox, K., Peherstorfer, B.: Learning nonlinear reduced models from data with operator inference. Annual Review of Fluid Mechanics **56**, 521–551 (2024)

[6] Engl, H.W., Hanke, M., Neubauer, A.: Regularization of Inverse Problems. Springer, Dordrecht (1996)

[7] Kolmogorov, A.N.: The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers. Doklady Akademii Nauk SSSR **30**, 301–305 (1941)

[8] Frisch, U.: Turbulence: The Legacy of A.N. Kolmogorov. Cambridge University Press, Cambridge, UK (1995)

[9] Chen, T., Chen, H.: Universal approximation to nonlinear operators by neural networks with arbitrary activation functions and its application to dynamical systems. IEEE Transactions on Neural Networks **6**(4), 911–917 (1995)

[10] Lu, L., Jin, P., Pang, G., Zhang, Z., Karniadakis, G.E.: Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. Nature Machine Intelligence **3**, 218–229 (2021)

[11] Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., Anandkumar, A.: Fourier neural operator for parametric partial differential equations. arXiv preprint arXiv:2010.08895 (2020)

[12] Kovachki, N., Li, Z., Liu, B., Azizzadenesheli, K., Bhattacharya, K., Stuart, A., Anandkumar, A.: Neural operator: Learning maps between function spaces. Journal of Machine Learning Research **24**(89), 1–97 (2023)

[13] Li, Z., Zheng, H., Kovachki, N., Jin, D., Chen, H., Liu-Schiaffini, M., Anandkumar, A.: Physics-informed neural operator for learning partial differential equations. arXiv preprint arXiv:2111.03794 (2021)

[14] Li, Z., Huang, D.Z., Liu-Schiaffini, M., Kovachki, N., Azizzadenesheli, K., Anandkumar, A.: Fourier neural operator with learned deformations for PDEs on general geometries. Journal of Machine Learning Research **24**(388), 1–26 (2023)

[15] Rahman, M.H., Ross, Z.E.: U-NO: U-shaped neural operators. arXiv preprint arXiv:2204.11127 (2022)

[16] Hao, Z., Su, C., Liu, S., Berner, J., Anandkumar, A., et al.: GNOT: A general neural operator transformer for operator learning. arXiv preprint arXiv:2302.14376 (2023)

[17] Wang, S., Perdikaris, P.: Improved architectures and training algorithms for deep operator networks. Journal of Scientific Computing **92**(2), 35 (2022)

[18] Ling, J., Kurzawski, A., Templeton, J.: Reynolds averaged turbulence modelling using deep neural networks with embedded invariance. Journal of Fluid Mechanics **807**, 155–166 (2016)

[19] Maulik, R., San, O.: Subgrid modelling for two-dimensional turbulence using neural networks. Journal of Fluid Mechanics **858**, 122–144 (2019)

[20] Loiseau, J.-C., Brunton, S.L.: Constrained sparse Galerkin regression. Journal of Fluid Mechanics **838**, 42–67 (2018)

[21] Kochkov, D., Smith, J.A., Alieva, A., Wang, Q., Brenner, M.P., Hoyer, S.: Machine learning–accelerated computational fluid dynamics. Proceedings of the National Academy of Sciences **118**(21) (2021)

[22] Fukami, K., Fukagata, K., Taira, K.: Super-resolution reconstruction of turbulent flows with machine learning. Journal of Fluid Mechanics **870**, 106–120 (2019)

[23] Fukami, K., Fukagata, K., Taira, K.: Machine-learning-based spatio-temporal super resolution reconstruction of turbulent flows. Journal of Fluid Mechanics **909**, 9 (2021)

[24] Liu, B., Tang, J., Huang, H., Lu, X.-Y.: Deep learning methods for super-resolution reconstruction of turbulent flows. Physics of Fluids **32**(2), 025105 (2020)

[25] Fukami, K., Fukagata, K., Taira, K.: Machine-learning-assisted flow reconstruction and super-resolution: a perspective. Annual Review of Fluid Mechanics (2023)

[26] Fukami, K., Fukagata, K., Taira, K.: Super-resolution analysis via machine learning: a survey for fluid flows. Theoretical and Computational Fluid Dynamics **37**(4), 421–444 (2023)

[27] Taira, K., Rigas, G., Fukami, K.: Machine learning in fluid dynamics: A critical assessment. Physical Review Fluids **10**(9), 090701 (2025)

[28] Dong, C., Loy, C.C., He, K., Tang, X.: Image super-resolution using deep convolutional networks. IEEE Transactions on Pattern Analysis and Machine Intelligence **38**(2), 295–307 (2016)

[29] Lim, B., Son, S., Kim, H., Nah, S., Lee, K.M.: Enhanced deep residual networks for single image super-resolution. In: CVPR Workshops, pp. 1132–1140 (2017)

[30] Stengel, K., Glaws, A., Hettinger, D., King, R.N.: Adversarial super-resolution of climatological wind and solar data. Proceedings of the National Academy of Sciences **117**(29), 16805–16815 (2020)

[31] Kraichnan, R.H.: Inertial ranges in two-dimensional turbulence. The Physics of Fluids **10**(7), 1417–1423 (1967)

[32] Boffetta, G., Ecke, R.E.: Two-dimensional turbulence. Annual Review of Fluid Mechanics **44**, 427–451 (2012)

[33] Taira, K., Brunton, S.L., Dawson, S.T.M., Rowley, C.W., Colonius, T., McKeon, B., Schmidt, O., Gordeyev, S., Theofilis, V., Ukeiley, L.S.: Modal analysis of fluid flows: An overview. AIAA Journal **55**(12), 4013–4041 (2017)

[34] Bao, K., *et al.*: Deep learning method for super-resolution reconstruction of turbulent flows: a review or perspective. Journal of Theoretical and Applied Mechanics Letters **13**, 100451 (2023)

[35] Lai, W.-S., Huang, J.-B., Ahuja, N., Yang, M.-H.: Deep Laplacian pyramid

networks for fast and accurate super-resolution. In: CVPR, pp. 624–632 (2017)

[36] He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR, pp. 770–778 (2016)

[37] Trottenberg, U., Oosterlee, C.W., Schüller, A.: Multigrid. Academic Press, San Diego, CA (2000)

[38] Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of GANs for improved quality, stability, and variation. In: International Conference on Learning Representations (ICLR) (2018)

[39] Bengio, Y., Louradour, J., Collobert, R., Weston, J.: Curriculum learning. In: Proceedings of the 26th International Conference on Machine Learning (ICML), pp. 41–48 (2009)

[40] Leith, C.E.: Diffusion approximation for two-dimensional turbulence. The Physics of Fluids **11**(3), 671–672 (1968)

[41] San, O., Staples, A.E.: High-order methods for decaying two-dimensional homogeneous isotropic turbulence. Computers & Fluids **63**, 105–127 (2012)

[42] Batchelor, G.K.: Computation of the energy spectrum in homogeneous two-dimensional turbulence. The Physics of Fluids Supplement II **12**, 233–239 (1969)

[43] Pathak, J., Subramanian, S., Harrington, P., Raja, S., Chattopadhyay, A., Mardani, M., Kurth, T., Hall, D., Li, Z., Azizzadenesheli, K., Hassanzadeh, P., Kashinath, K., Anandkumar, A.: FourCastNet: A global data-driven high-resolution weather model using adaptive fourier neural operators. arXiv preprint arXiv:2202.11214 (2022)

[44] Boullé, N., Halikias, G., Otto, F., Townsend, A.: Operator learning without the adjoint. Journal of Machine Learning Research **25**, 1–86 (2024)

[45] Kovachki, N.B., Lanthaler, S., Stuart, A.M.: Operator learning: Algorithms and analysis. arXiv preprint arXiv:2402.15715 (2024)

[46] Maulik, R., San, O., Rasheed, A., Vedula, P.: Subgrid modelling for two-dimensional turbulence using neural networks: approaches and lessons. Theoretical and Computational Fluid Dynamics **34**, 251–273 (2020)

[47] Beck, A., Flad, D., Munz, C.-D.: Deep neural networks for data-driven LES closure models. Journal of Computational Physics **398**, 108910 (2019)

[48] Wu, J.-L., Xiao, H., Paterson, E.: Physics-informed machine learning approach for augmenting turbulence models: a comprehensive framework. Physical Review Fluids **3**(7), 074602 (2018)

[49] Fukami, K., Fukagata, K.: Single-snapshot machine learning for super-resolution of turbulence. Journal of Fluid Mechanics **998**, 17 (2024)

[50] Fukami, K., Taira, K.: Single-snapshot machine learning for super-resolution of turbulence. Journal of Fluid Mechanics **1001**, 32 (2024)

[51] Fukami, K., Taira, K.: Observable-augmented manifold learning for multi-source turbulent flow data. Journal of Fluid Mechanics **1010**, 4 (2025)

[52] Gao, H., Sun, L., Wang, J.-X.: Super-resolution and denoising of fluid flow using physics-informed convolutional neural networks without high-resolution labels. Physics of Fluids **33**(7), 073603 (2021)

[53] Kim, H., *et al.*: Unsupervised deep learning for super-resolution reconstruction of turbulence. Journal of Fluid Mechanics **910**, 29 (2021)

[54] Sanchez-Gonzalez, A., Godwin, J., Pfaff, T., Ying, R., Leskovec, J., Battaglia, P.: Learning to simulate complex physics with graph networks. In: International Conference on Machine Learning (ICML), pp. 8459–8468 (2020)

[55] Ozawa, Y., Honda, H., Nonomura, T.: Spatial superresolution based on simultaneous dual PIV measurement with different magnification. Experiments in Fluids **65**(4), 42 (2024)

[56] Stahl, S.L., Benton, S.I.: Sparse flow reconstruction methods to reduce the costs of analyzing large unsteady datasets. Journal of Computational Physics **534**, 114037 (2025)

[57] Yousif, M.Z., Yu, L., Lim, H.-C.: A deep-learning approach for reconstructing 3D turbulent flows from 2D observation data. Scientific Reports **13**, 29525 (2023)

[58] Pang, Z., *et al.*: A deep-learning super-resolution reconstruction model of subgrid-scale turbulent flow fields in large-eddy simulation. Computers & Structures **298**, 107392 (2024)

[59] Yu, L., Chen, Y., Yousif, M.Z., Lim, H.-C.: Three-dimensional super-resolution reconstruction of turbulent flow using 3D-ESRGAN with random sampling strategy. Computers & Fluids, 106890 (2025)

[60] Wu, H.-L., Xu, A., Xi, H.-D.: Super-resolution reconstruction of turbulent flows from a single Lagrangian trajectory. Journal of Fluid Mechanics **1026**, 46 (2026)

[61] Wu, H., Cao, Y., Chen, Y., Laima, S., Chen, W.-L., Zhou, D., Li, H.: Generalizable super-resolution turbulence reconstruction from minimal training data. Journal of Fluid Mechanics **1024**, 27 (2025)

[62] Ye, M.-k., Fan, G.-q., Wan, D.-c.: Enhancing turbulent channel flow super-resolution via physics-guided deep learning with global pressure correlations.

Journal of Hydrodynamics, 1–10 (2025)

[63] Page, J., *et al.*: Super-resolution of turbulence with dynamics in the loss. Journal of Fluid Mechanics **1011**, 38 (2025)

[64] Page, J.: Super-resolution of turbulence with dynamics in the loss. Journal of Fluid Mechanics **1002**, 3 (2025)

[65] Price, I., Sanchez-Gonzalez, A., Alet, F., Ewalds, T., *et al.*: Probabilistic weather forecasting with machine learning. Nature **637**, 84–90 (2024)

[66] Kim, J., Lee, J.K., Lee, K.M.: Accurate image super-resolution using very deep convolutional networks. In: CVPR, pp. 1646–1654 (2016)

[67] Kim, J., Lee, J.K., Lee, K.M.: Deeply-recursive convolutional network for image super-resolution. In: CVPR, pp. 1637–1645 (2016)

[68] Tai, Y., Yang, J., Liu, X.: Image super-resolution via deep recursive residual network. In: CVPR, pp. 3147–3155 (2017)

[69] Dong, C., Loy, C.C., Tang, X.: Accelerating the super-resolution convolutional neural network. In: European Conference on Computer Vision, pp. 391–407 (2016)

[70] Shi, W., Caballero, J., Huszár, F., Totz, J., Aitken, A.P., Bishop, R., Rueckert, D., Wang, Z.: Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In: CVPR, pp. 1874–1883 (2016)

[71] Zhang, Y., Tian, Y., Kong, Y., Zhong, B., Fu, Y.: Residual dense network for image super-resolution. In: CVPR, pp. 2472–2481 (2018)

[72] Haris, M., Shakhnarovich, G., Ukita, N.: Deep back-projection networks for super-resolution. In: CVPR, pp. 1664–1673 (2018)

[73] Zhang, Y., Li, K., Li, K., Wang, L., Zhong, B., Fu, Y.: Image super-resolution using very deep residual channel attention networks. In: ECCV, pp. 286–301 (2018)

[74] Wang, X., Yu, K., Wu, S., Gu, J., Liu, Y., Dong, C., Qiao, Y., Change Loy, C.: ESRGAN: Enhanced super-resolution generative adversarial networks. In: ECCV Workshops (2018)

[75] Liang, J., Cao, J., Sun, G., Zhang, K., Van Gool, L., Timofte, R.: SwinIR: Image restoration using Swin transformer. In: ICCV Workshops (2021)

[76] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. Advances in Neural Information Processing Systems **30** (2017)

[77] Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D.J., Norouzi, M.: Image super-resolution via iterative refinement. In: IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 45, pp. 4713–4726 (2022)

[78] Ahn, N., Kang, B., Sohn, K.-A.: Fast, accurate, and lightweight super-resolution with cascading residual network. In: ECCV, pp. 252–268 (2018)

[79] Kong, X., Zhao, H., Qiao, Y., Dong, C.: ClassSR: A general framework to accelerate super-resolution networks by data characteristic. In: CVPR, pp. 12016–12025 (2021)

[80] Chandler, G.J., Kerswell, R.R.: Invariant recurrent solutions embedded in a turbulent two-dimensional Kolmogorov flow. Journal of Fluid Mechanics **722**, 554–595 (2013)

[81] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: PyTorch: An imperative style, high-performance deep learning library. In: Advances in Neural Information Processing Systems vol. 32, (2019)

[82] Mallat, S.G.: A theory for multiresolution signal decomposition: The wavelet representation. IEEE Transactions on Pattern Analysis and Machine Intelligence **11**(7), 674–693 (1989)

[83] Canuto, C., Hussaini, M.Y., Quarteroni, A., Zang, T.A.: Spectral Methods in Fluid Dynamics. Springer Series in Computational Physics. Springer, New York (1988)

[84] Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)