

State estimation in homogeneous isotropic turbulence using super-resolution with a 4DVar training algorithm

Markus Weyrauch^{1,2} , Moritz Linkmann²  and Jacob Page² 

¹Institute for Water and Environment – Numerical Fluid Mechanics Group, Karlsruhe Institute of Technology, 76131 Karlsruhe, Germany

²School of Mathematics and Maxwell Institute for Mathematical Sciences, University of Edinburgh, Edinburgh EH9 3FD, UK

Corresponding author: Markus Weyrauch, markus.scherer@kit.edu

(Received 20 October 2025; revised 26 February 2026; accepted 26 February 2026)

Variational data assimilation and machine-learning based super-resolution are two alternative approaches to state estimation in turbulent flows. The former is an optimisation problem featuring a time series of coarse observations, the latter usually requires a library of high-resolution ‘ground truth’ data. We show that the classic ‘4DVar’ data assimilation algorithm can be used to train neural networks for super-resolution in three-dimensional isotropic turbulence without the need for high-resolution reference data. To do this, we adapt a pseudo-spectral version of the fully differentiable JAX-CFD solver (Kochkov *et al.*, *Proc. Natl Acad. Sci. USA*, vol. 118, issue 21, 2021, e2101784118) to three-dimensional flows and combine it with a convolutional neural network for super-resolution. As a result, we are able to include entire trajectories in our loss function, which is minimised with gradient-based optimisation to define the neural network weights. We show that the resulting neural networks outperform 4DVar for state estimation at initial time over a wide variety of metrics, though 4DVar leads to more robust predictions towards the end of its assimilation window. We also present a hybrid approach in which the trained neural network output is used to initialise 4DVar. The resulting performance is more than twice as accurate as other state estimation strategies for all times and performs well even beyond known limiting length scales, all without requiring access to high-resolution measurements at any point.

Key words: homogeneous turbulence, machine learning, low-dimensional models

1. Introduction

Chaotic dynamical systems are notoriously sensitive to small variations in the initial conditions. Numerical forecasts initiated with an imperfect representation of the current system state will therefore deviate from the evolution of the true system, regardless of how accurately the numerical model captures the governing physics. A popular family of methods to account for these deviations in numerical forecasting techniques is data assimilation, which augments the numerical prediction in regular intervals with observations of the true system state to ensure that the prediction stays in the neighbourhood of the true system trajectory (Talagrand 1997). These state estimation techniques usually rely on comparison of the forward propagation of predictions in time, with comparison to observational data. They are complemented by recent alternatives to the state estimation problem that utilise deep neural networks. This approach, usually termed ‘super-resolution’ (see, e.g. Fukami, Fukagata & Taira 2023), can perform well with limited observations, but usually relies on the existence of a large dataset of high-resolution snapshots for network training. This paper seeks to combine both approaches in a three-dimensional turbulent flow, infusing the training process of a neural network with a data assimilation algorithm (which does not rely on high-resolution reference data) and initialising assimilation computations with the output of a neural network.

The current state-of-the-art in data assimilation are variational methods, ensemble Kalman filters or combinations of both (Bannister 2017). Classical Kalman filters (Kalman 1960) are sequential techniques that first advance a system state in time until observations become available. New measurements are used to improve the prediction, with individual observations weighted depending on their uncertainty (Evensen 2007). For large-scale nonlinear systems, ensemble Kalman filters replace a single trajectory with an ensemble of initial conditions and update the prediction based on statistics of the ensemble (Evensen 1994).

Variational techniques like 4DVar (where 4D denotes three spatial dimensions and time) solve a constrained optimisation problem to determine the initial condition that minimises the deviation between observations and predictions over an assimilation window $0 \leq t \leq T_{DA}$ in a least squares sense (Talagrand & Courtier 1987). Gradients of the predictive model with respect to the initial condition are usually computed via adjoint methods (Kalnay 2002). An alternative to coding and backward-time marching an adjoint system is the use of a fully differentiable solver. In this approach, gradients are computed via automatic differentiation through the computational graph. Automatic differentiation back-propagates gradients through the forward computation using a combination of the chain rule (which requires storage of the forward state at each node in the graph) with known derivatives of elementary mathematical operations (Baydin *et al.* 2018). As a result, gradients are accurate to machine precision.

In the past, both 4DVar (Bewley & Protas 2004; Gronskis, Heitz & Mémin 2013; Wang, Wang & Zaki 2019; Li *et al.* 2020; Wang, Wang & Zaki 2022) and (ensemble) Kalman filters (Höpfner *et al.* 2005; Chevalier *et al.* 2006; Colburn, Cessna & Bewley 2011; Kato & Obayashi 2013; Kato *et al.* 2015) have been applied in a variety of transitional and fully developed turbulent flows, with a strong focus on wall-bounded configurations (Wang *et al.* 2019; Buchta & Zaki 2021; Wang & Zaki 2021; Wang *et al.* 2022; Wang & Zaki 2025); see Hayase (2015), Mons, Chassaing & Sagaut (2017) and Zaki (2025) for a comprehensive overview. We consider here the state estimation problem in triply periodic, body-forced turbulence. To our knowledge, 4DVar has been considered in this configuration

only by Li *et al.* (2020), who drove the flow with a monochromatic, unidirectional body force.

In Li *et al.* (2020) the observations available to the 4DVar algorithm are low-pass filtered velocity snapshots from a target trajectory. State estimation is robust when the cutoff wavenumber of the spectral filter is not significantly smaller than a critical value $k_C \approx 0.2\eta_K^{-1}$ (equivalent to a critical length scale $l_C \approx 5\pi\eta_K$, with η_K denoting the Kolmogorov length). The fact that 4DVar struggles as the observational data are coarsened in space is consistent with earlier findings on turbulence synchronisation. Numerical studies of chaos synchronisation in turbulent flows have found that the small scales (wavenumbers $k > k_C$) of a turbulent flow can be reconstructed by time-stepping high-pass-filtered Navier–Stokes equations if all modes up to a wavenumber $k \leq k_C$ are prescribed (Yoshida, Yamaguchi & Kaneda 2005; Lalescu, Meneveau & Eyink 2013; Vela-Martín 2021; Inubushi & Caulfield 2026). The indication is that the small-scale dynamics, $k > k_C$, is slaved to the large-scale evolution. Note that a slightly different problem has also been considered in a channel (Wang & Zaki 2022), where synchronisation of masked vertical slabs occurs if the height of the masked region is less than the local Taylor microscale.

With the increasing popularity of machine-learning techniques (Brenner, Eldredge & Freund 2019; Brunton, Noack & Koumoutsakos 2020; Taira, Rigas & Fukami 2025), a variety of data-driven alternatives for state estimation have emerged, though connections with classical state estimation methods are not usually emphasised. Popular examples are super-resolution methods that were originally developed for image processing. In these approaches, a model, typically a convolutional neural network (CNN), is trained to reconstruct high-resolution images based on low-quality inputs. This is accomplished by training the model to minimise a reconstruction error on a library of high-resolution target images (Dong *et al.* 2016). These methods are straightforward to adapt to turbulent flows, where the goal is the reconstruction of high-resolution flow fields from coarse observations. The networks learn data-driven interpolation operators whose ability to reconstruct high-resolution fields outperforms standard polynomial interpolation (Fukami *et al.* 2019, 2021; Kim *et al.* 2021). Recently, more advanced super-resolution techniques have incorporated additional terms in the training loss functions to encourage the reconstructions to satisfy physical constraints (e.g. mass conservation) or to improve reconstruction of velocity gradients/vorticity; a detailed review of different approaches and model architectures can be found in Fukami *et al.* (2023).

We view super-resolution and data assimilation techniques as complementary approaches to state estimation: while 4DVar estimates the state given observations of a single trajectory, super-resolution models see a wide variety of flow states sampled from the turbulent attractor but have no knowledge of the dynamics along any specific trajectory. Recent work in data assimilation hints at the potential for improving classical assimilation techniques with the incorporation of super-resolution methods. Some examples include neural networks that increase the flow field resolution before the assimilation step in ensemble Kalman filters (Barthélémy *et al.* 2022) or techniques in which data assimilation and super-resolution are performed simultaneously by a neural network (Yasuda & Onishi 2023). Another promising approach is to generate high-quality initial guesses for 4DVar with the aid of super-resolution as in Frerix *et al.* (2021), where a CNN learned an approximate inverse of the observation operator that maps physical data onto the observation space.

A caveat most of these hybrid techniques share with pure super-resolution methods is a continued reliance on high-resolution training data, which limits their applicability

in situations where large, fully resolved datasets for training are not available (e.g. experimental measurements). Interestingly, recent observations indicate that models trained exclusively on low-resolution data can estimate high-resolution fields when physical principles are incorporated in the training process. For example, Kelshaw, Rigas & Magri (2022) trained such a ‘physics-informed’ network on low Reynolds number (Re) two-dimensional Kolmogorov flow with a loss function that includes a penalisation when the network output does not satisfy the Navier–Stokes equations at each node of the coarse grid. The same group have also recently presented a method to generate three-dimensional flow fields from planar observations in a low- Re , chaotic Kolmogorov flow (Mo & Magri 2025). In another approach, Page (2025*b*) (hereafter P25) designed a loss function in analogy to the classical 4DVar optimisation problem by including a numerical time marching scheme in the training process (‘solver-in-the-loop’, Um *et al.* 2020). In contrast to standard super-resolution training on individual high-resolution fields, the super-resolved low-resolution snapshots are time marched using a fully differentiable direct numerical simulation (DNS) code. This trajectory is then coarse-grained for comparison with the reference coarse observations, just as in classical 4DVar. But while 4DVar seeks the high-resolution initial condition that minimises the deviation between observations and predictions over the assimilation window, trajectory-based super-resolution models learn how to reconstruct such an optimal initial condition from a coarse-grained observation of it. Compared with standard super-resolution, the trajectory-based approach hence compensates for its lack of access to high-resolution reference data by learning how to reconstruct high-resolution trajectories that faithfully reproduce the target trajectory over a given time window. Page (2025*b*) demonstrated the method in two-dimensional Kolmogorov turbulence at fairly high Re , where performance of the model was found to be comparable to high-resolution-based super-resolution approaches. It also outperforms classical 4DVar at the initial state estimation for all the levels of coarse-graining that were tested, though the model predictions deviate much faster from the reference trajectory than the 4DVar approach once unrolled in time.

In this paper we consider the state estimation problem in three-dimensional, homogeneous, isotropic turbulence, using combinations of neural network super-resolution and 4DVar – none of which require reference high-resolution data. We first adapt the 4DVar-inspired super-resolution approach of P25 to the three-dimensional problem and assess its capabilities. To do this, we have developed a spectral, fully differentiable three-dimensional Navier–Stokes solver around the JAX-CFD data structures introduced in the landmark work by Kochkov *et al.* (2021) and Dresdner *et al.* (2023). This allows us to include entire trajectories of the three-dimensional flow in our loss function for network training, and also to perform classical 4DVar using automatic differentiation. We explore the limits of the super-resolution based approach with respect to known limiting length scales, and also demonstrate that 4DVar calculations initialised with super-resolved outputs perform more than two times better than other methods for state reconstruction over the entire assimilation window and beyond.

The remainder of this paper is organised as follows: In § 2 the flow configuration is presented and relevant physical measures are introduced. The essentials of the differentiable DNS solver are summarised in § 3, alongside an overview of the data assimilation and super-resolution techniques used in this study. In § 4 the predictive performance of our trajectory-based super-resolution technique is compared with other state estimation methods and the potential of using the super-resolved fields for the initialisation of 4DVar is assessed. The paper closes with a summary of the relevant findings and an outlook on future work in § 5.

2. Flow configuration

We consider the non-dimensional Navier–Stokes equations for an incompressible Newtonian fluid under the action of an external body force \mathbf{f}_e ,

$$\partial_t \mathbf{u} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\nabla p + \frac{1}{Re} \Delta \mathbf{u} + \mathbf{f}_e, \quad (2.1)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (2.2)$$

in a triply periodic domain $\Omega = [0, L]^3$. Velocity, vorticity and pressure fields are denoted as $\mathbf{u} = (u, v, w)^T$, $\boldsymbol{\omega} = \nabla \times \mathbf{u} = (\omega_x, \omega_y, \omega_z)^T$ and p , respectively.

Equation (2.1) has been non-dimensionalised by the characteristic length and time scales $L/2\pi$ and $\sqrt{L/(2\pi\chi)}$, respectively, leading to a Reynolds number $Re = (\sqrt{\chi}/\nu)(L/2\pi)^{3/2}$. Here, ν and χ denote the kinematic viscosity and a representative forcing amplitude per unit mass (forcing profile defined below), respectively. A deterministic, isotropic and non-helical forcing is chosen to sustain turbulent fluctuations by continuously injecting energy into a range of low wavenumbers (McKay *et al.* 2017; Linkmann 2018). In dimensional form, the forcing is

$$\mathbf{f}_e^* = \chi \sum_{k=1}^{k_f} \begin{pmatrix} \sin(2\pi k z^*/L) + \sin(2\pi k y^*/L) \\ \sin(2\pi k x^*/L) + \sin(2\pi k z^*/L) \\ \sin(2\pi k y^*/L) + \sin(2\pi k x^*/L) \end{pmatrix}, \quad (2.3)$$

with the highest forcing wavenumber set at $k_f = 3$. Note that the forcing amplitude, χ , is used to define the time scale (and hence Re) in the problem, and so is not specified directly.

In order to assess the ability of the various state estimation techniques to accurately reproduce turbulent snapshots, we compare various spectral quantities that we introduce here for convenience. We focus on standard observables assessing the energy content at a given scale, that is, the kinetic energy spectrum

$$E(k, t) = \frac{1}{2} \int_{|\mathbf{k}|=k} d\mathbf{k} |\hat{\mathbf{u}}(\mathbf{k}, t)|^2, \quad (2.4)$$

where $\hat{\mathbf{u}}$ is the Fourier transform of the velocity field, and the mean flux of kinetic energy across scales

$$\Pi(k, t) = - \left\langle \tilde{S}_{ij}^k \left(u_i \widetilde{u}_j^k - \widetilde{u}_i^k \widetilde{u}_j^k \right) \right\rangle_{\Omega}, \quad (2.5)$$

where $(\widetilde{\cdot})^k$ refers to Galerkin truncation at wavenumber k and $S_{ij} = (\partial_i u_j + \partial_j u_i)/2$ is the rate-of-strain tensor. The angled brackets denote a spatial average over the domain Ω . We also use the second and third invariants of the velocity gradient tensor (Q and R , Chong, Perry & Cantwell 1990) to assess our small-scale reconstructions.

Throughout this paper, time and length scales will be measured in terms of the large-eddy time $T_e = u_{rms}^2/\bar{\varepsilon}$ and the Kolmogorov length $\eta_K = (\nu^3/\bar{\varepsilon})^{1/4}$, respectively, where $u_{rms} = \sqrt{2\bar{E}/3}$ is the root-mean-square velocity calculated from the spatio-temporally averaged kinetic energy, \bar{E} , in a statistically steady state and $\bar{\varepsilon}$ is the spatio-temporal average of the dissipation rate $\varepsilon(\mathbf{x}, t) = 2\nu S_{ij} S_{ji}$. The Taylor and large-eddy Reynolds numbers are defined as $Re_\lambda = u_{rms} \lambda_T/\nu$ and $Re_L = \bar{E}^{1/2} L_e/\nu = (3/20) Re_\lambda^2$, respectively, where $\lambda_T = \sqrt{15 u_{rms}^2 \nu/\bar{\varepsilon}}$ is the Taylor microscale and $L_e = (\bar{E})^{3/2}/\bar{\varepsilon}$ is a length scale associated with the large eddies (Pope 2000).

Table 1 summarises the relevant physical and numerical parameters of the simulations that will be analysed in the remainder of this study. The values are obtained from ensemble

Setting	N_i	$k_{(m,N)}\eta_K$	$\Delta x_{(c)}/\eta_K$	$\Delta x_{(c)}/l_C$	Re_L	Re_λ	L/η_K	L_e/η_K	λ_T/η_K
DNS	128	1.52	1.36	0.09	689.04	67.62	174.02	134.30	6.18
$M=4$	32	0.58	5.44	0.35					
$M=8$	16	0.29	10.88	0.69					
$M=16$	8	0.14	21.75	1.38					

Table 1. Physical and numerical parameters of the DNS database used in this study, together with information on the low-resolution fields (coarsening factor $\equiv M$). Here N_i denotes the number of grid nodes per spatial direction and k_m is the largest resolved wavenumber, which we scale here with the Kolmogorov length η_K . For the coarsened fields, we report the Nyquist wavenumber k_N scaled by η_K . Additionally, Δx (DNS run) and Δx_c (coarsened fields) are the grid spacings in the DNS ground truth and the coarsened fields, respectively, while l_C is the critical length scale for turbulence synchronisation. For the full DNS, we also provide values for the large-eddy (Re_L) and Taylor Reynolds numbers (Re_λ). The last three columns summarise some key length scale ratios. All super-resolution models are trained with an unroll time $T_u = 0.22T_e$, while 4DVar-based reconstructions are shown for three different assimilation windows of lengths $T_{DA}/T_e \in \{0.4, 0.9, 1.3\}$, respectively, where T_e is the large-eddy time.

averages over 100 independent trajectories each of length $88T_e$ (equivalent to 100 non-dimensional time units). At the chosen parameter point, the flow attains a Taylor Reynolds number of $Re_\lambda \approx 70$ (equivalent to $Re = 50$ and $Re_L \approx 690$), leading to a separation between the system’s largest and smallest scales of $L_e/\eta_K \approx 134$. These values are comparable to Li *et al.* (2020) who had $Re_\lambda \approx 75$, albeit with a different forcing profile (see Appendix B for an assessment of 4DVar in their flow). We point out that even though an inertial range cannot be identified at this Reynolds number, velocity gradient statistics are non-Gaussian, indicating the presence of small-scale intermittency (Schumacher *et al.* 2014) – in other words, velocity increment statistics are not self-similar across different scales: a key feature of fully developed three-dimensional turbulence is present in the data considered here.

3. Numerical method

We describe here the differentiable DNS solver used in this study, along with the candidate state estimation procedures that are the focus of this paper. These include the classic 4DVar algorithm and the machine learning ‘super-resolution’ approach inspired by variational data assimilation.

3.1. The DNS code

The pseudo-spectral DNS solver used in this study is an adaptation of the two-dimensional, spectral version of the JAX-CFD solver (Dresdner *et al.* 2023). We implement the three-dimensional problem using the velocity–vorticity formulation proposed by Kim, Moin & Moser (1987), maintaining the underlying GPU efficiency of the JAX-CFD time-stepping routines and associated data structures (Kochkov *et al.* 2021; Dresdner *et al.* 2023), which was built around the JAX library (Bradbury *et al.* 2018). The use of the JAX library means that gradients of functions with respect to their inputs can be computed in a single line of code via back-propagation – values are stored in the forward pass through the compute graph to allow for application of the chain rule (‘automatic differentiation’) combined with known primitives. The differentiability of the two-dimensional version of the code has previously been used to find unstable periodic orbits (Page *et al.* 2024; Page 2025a) and for the design of mixing strategies in complex fluids (Alhashim, Hausknecht & Brenner 2025). This flexibility means that it is straightforward to consider other assimilation-like




Figure 1. Sample slices of the vorticity field ω_y/ω_{rms} at $z=0$. Black and red lines indicate the Taylor microscale λ_T and the critical length scale $l_C = 5\pi\eta_K$, respectively. The white grid lines in the lower left corners of the first three frames visualise the coarsened grids for coarsening factors $M \in \{4, 8, 16\}$ relative to the DNS grid.

problems, e.g. optimising for an unknown forcing, with no alteration to the underlying solver.

The governing equations in the velocity–vorticity formulation of Kim *et al.* (1987) are

$$\partial_t \mathbf{q} = \mathbf{h} + \frac{1}{Re} \Delta \mathbf{q}, \tag{3.1}$$

where $\mathbf{q} := (\Delta v, \omega_y)^T$ and $\mathbf{h} := (h_v, h_\omega)^T$ with

$$h_v = -\partial_y (\partial_x H_1 + \partial_z H_3) + (\partial_x^2 + \partial_z^2) H_2, \tag{3.2}$$

$$h_\omega = \partial_z H_1 - \partial_x H_3. \tag{3.3}$$

In (3.2) and (3.3) the term $\mathbf{H} = (H_1, H_2, H_3)^T := -(\mathbf{u} \cdot \nabla)\mathbf{u} + \mathbf{f}_e$ combines the nonlinear convective term and the external body force \mathbf{f}_e . Both \mathbf{u} and $\boldsymbol{\omega}$ are expanded as truncated Fourier series in all three spatial dimensions and are dealiased according to the 2/3 rule. The nonlinear (\mathbf{H}) terms are advanced using a fourth-order Runge Kutta method, while the diffusion term is treated with a semi-implicit Crank–Nicholson (CN) scheme. The use of the CN scheme, rather than an integrating factor as usually done in the standard Fourier-based pseudo-spectral method, was done simply for ease of interaction with JAX-CFD time-stepping data structures. A validation of the differentiable DNS code is provided in Appendix A.

In all simulations, the velocity field is discretised by $N_x \times N_y \times N_z = 128^3$ collocation points, with a maximum resolved wavenumber $k_m \eta_K = 1.52$. Three different coarse-graining factors $M \in \{4, 8, 16\}$ between the fine DNS grid and the coarser observation grids are considered for the state estimation problem – while the lower two values fall below the synchronisation limit $l_C = 5\pi\eta_K \approx 11.55\Delta x$ (Lalescu *et al.* 2013), the most severe coarsening leads to a grid width about 40% larger than l_C (see also table 1). In figure 1 the coarsened grids are visualised in front of representative high-resolution snapshots of the vorticity, together with the Taylor microscale λ_T and the synchronisation limit l_C . At $M = 4$, the mesh is sufficiently fine to resolve the vortical structures visible in the figure, while $M = 16$ is visibly much wider than the individual flow structures.

3.2. Data assimilation

Considering the Navier–Stokes equations as a perfect forecasting model, a classical 4DVar approach seeks an initial flow state $\mathbf{v}_0 \in \mathcal{P}$ in the physical ‘target’ space \mathcal{P} that minimises a loss function

$$\mathcal{L}_{DA}(\mathbf{v}_0) = \frac{1}{N_T} \sum_{k=0}^{N_T-1} \|\mathbf{m}_{t_k} - \mathcal{H} \circ \boldsymbol{\varphi}_{t_k}(\mathbf{v}_0)\|^2, \tag{3.4}$$

where $\varphi_t(\mathbf{u})$ is the time forward map associated with the governing equations (which in our formulation includes back-and-forth conversion from primitive variables to the velocity–vorticity form), the variable $\mathbf{m}_{t_k} \in \mathcal{O}$ corresponds to one of N_T measurements/observations (\mathcal{O} is an abstract space of observations) and $\mathcal{H} : \mathcal{P} \rightarrow \mathcal{O}$ maps the velocity field into observations. Assimilation is performed over a time window $t_k \in [0, T_{DA}]$, and the objective of minimising (3.4) corresponds to finding an initial condition \mathbf{v}_0 that best reproduces the time series of measurements as it is unrolled in time.

In the current study, the observation operator (\mathcal{H} above) is a simple downsampling operation that we denote \mathcal{C} . The physical space $\mathcal{P} = \mathbb{R}^{N_x \times N_y \times N_z \times 3}$ includes our high-resolution fields, while the observation space $\mathcal{O} = \mathbb{R}^{N_x/M \times N_y/M \times N_z/M \times 3}$ contains the corresponding coarse-grained fields. The optimisation problem (3.4) therefore becomes

$$\mathcal{L}_{DA}(\mathbf{v}_0) = \frac{1}{N_T} \sum_{k=0}^{N_T-1} \|\mathbf{m}_{t_k} - \mathcal{C} \circ \varphi_{t_k}(\mathbf{v}_0)\|^2, \tag{3.5}$$

where the time series of observations or measurements $\{\mathbf{m}_{t_k}\}_{k=0}^{N_T-1}$ correspond here to coarse-grained ground truth snapshots $\mathbf{m}_{t_k} \equiv \mathcal{C} \circ \varphi_{t_k}(\mathbf{u}_0)$ obtained from an (unseen) DNS, but could be provided from another source (e.g. an experiment). We compute gradients of the loss (via automatic differentiation) with respect to the initial condition, $\nabla_{\mathbf{v}_0} \mathcal{L}_{DA}$, and supply them to an Adam optimiser (Kingma & Ba 2015) with initial learning rate $\eta = 0.1$. In all computations, we take 200 gradient update steps and select the best-performing \mathbf{v}_0 . The choice of the initial guess for the optimisation procedure will be discussed in the next section.

3.3. Super-resolution

Standard super-resolution approaches train a neural network to generate high-resolution images (here velocity fields from a high-fidelity DNS run) from low-resolution datasets. The model thus learns an interpolation operator $\mathcal{N}_\theta : \mathcal{O} \rightarrow \mathcal{P}$ (i.e. the CNN) as an approximate inverse to the coarse-graining operator \mathcal{C} . The variable θ indicates the network weights that minimise a loss of the form (Fukami *et al.* 2019, 2021):

$$\mathcal{L}_{SR} = \frac{1}{N_S} \sum_{j=1}^{N_S} \|\mathbf{u}_j - \mathcal{N}_\theta \circ \mathcal{C}(\mathbf{u}_j)\|^2. \tag{3.6}$$

Training is performed on N_S individual snapshots (here velocity fields \mathbf{u}_j). As discussed in § 1, training a neural network with the loss (3.6) can lead to ‘unphysical’ velocity fields, e.g. not divergence free. Therefore, additional terms or modifications to the loss are typically required to ensure that the reconstructed fields fulfil at least some of the desired physical properties (see Fukami *et al.* 2023 and references therein).

In analogy to the 4DVar loss function (3.5) above, the loss in our trajectory-based super-resolution technique (hereafter ‘SRdyn’) contains low-resolution trajectories that are obtained from unrolling the coarse-grained training dataset over a fixed unroll time T_u (Page 2025b):

$$\mathcal{L}_{TC} = \frac{1}{N_S N_T} \sum_{j=1}^{N_S} \sum_{k=0}^{N_T-1} \left\| \mathbf{m}_{t_k}^j - \mathcal{C} \circ \varphi_{t_k} \circ \mathcal{N}_\theta(\mathbf{m}_{t_0}^j) \right\|^2. \tag{3.7}$$

The loss (3.7) measures the deviation between a time series of measurements, $\{\mathbf{m}_{t_k}^j\}_{k=0}^{N_T-1}$, and the coarse-grained trajectory obtained by time marching the high-resolution fields produced by the neural network and coarsening the output. As in the discussion of 4DVar above, the measurements here are obtained by coarse-graining a reference DNS trajectory, $\mathbf{m}_{t_k}^j \equiv \mathcal{C} \circ \boldsymbol{\varphi}_{t_k}(\mathbf{u}_j)$, but the explicit ‘measurement’ notation is used to show that these observations could be obtained from a completely different source (an experiment, another simulation etc.). The deviation from the reference observations is evaluated at N_T discrete times $t_k \in \{0, \Delta t_c, \dots, (N_T - 1)\Delta t_c\}$ along each trajectory in the interval $[0, T_u]$, where $\Delta t_c = M_T \Delta t$ denotes the time interval between two consecutive observation times and Δt is the simulation time step in the DNS. We choose the temporal coarsening factor to match the degree of spatial coarse-graining, $M_T = M$, to mimic the limited temporal availability of observational data, for instance, in an experiment. The neural network includes a Leray projection $\mathcal{L}_u(\tilde{\mathbf{u}}) := \tilde{\mathbf{u}} - \nabla \Delta^{-1} \nabla \cdot \tilde{\mathbf{u}}$ as its final ‘layer’ to guarantee its output to be divergence free. The Leray projection is performed in Fourier space before inversion back to physical space (all super-resolution/assimilation is done in physical space). The inclusion of a Leray projection in the training loop is also possible in domains with non-periodic boundary conditions, though solution of the corresponding Poisson problem will be more costly.

While in this work the measurements from which the neural network seeks to construct a full state are generated by the same solver called in the training loop, the approach described above is general and allows for the use of measurements obtained in experiments or in other separate computations. The training algorithm seeks to combine positive aspects of 4DVar, i.e. matching a reference trajectory, and super-resolution which succeeds when the model is exposed to enough data to parameterise the solution manifold of the underlying differential equation. Once trained, the neural network takes a single measurement snapshot (here coarse-grained velocity) and predicts the full state. Using the solver in the loop, rather than as a mechanism to generate a large dataset for the classical super-resolution approach (3.6), has additional benefits that extend beyond the present study. For example, network training can immediately begin once measurements are available without generation of a reference dataset, which is helpful if there is uncertainty in parameters, boundary conditions, inflow/outflow conditions etc. Moreover, we believe the approach can be particularly promising in non-Newtonian problems where one may also want to adjust the underlying constitutive model as part of the optimisation.

All SRdyn models discussed here were trained to minimise the loss (3.7) using a database of $N_{traj} = 75$ independent trajectories, each of length $88T_e$ (equivalent to 100 non-dimensional time units), with 50 snapshots stored per trajectory (3750 total snapshots). Each trajectory has been initialised with an arbitrary turbulent flow state that was obtained from time stepping a random solenoidal velocity field until a statistically stationary state was reached. During the training process, the model sees the snapshots of the training dataset in random order, without considering whether or not they belong to the same original (long) trajectory. We remove 10 % of the training dataset for validation (to verify the model is not over-fitting). The results presented in § 4 are based on a separate ‘test’ set of velocity field snapshots obtained in the same way. Training is performed with an Adam optimiser (Kingma & Ba 2015) with a learning rate $\eta = 10^{-4}$. We train for between 25 and 50 epochs (complete cycles through the training data) with a batch size of 16 individual trajectories (note the ability to run many independent DNS calculations simultaneously on a GPU). We explored several unroll times in a range $T_u/T_e \in [0.22, 1.32]$, which indicated a rather weak impact of this parameter on the performance of the trained models. Given this and the fact that the time-marching of the DNS is the most intensive aspect of the

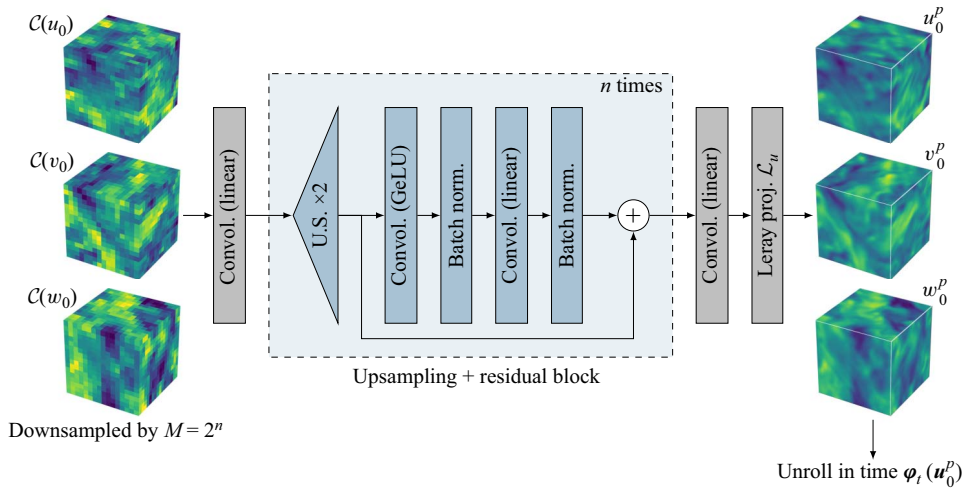


Figure 2. Neural network architecture used to train the SRdyn model with ‘U.S.’ indicating an upsampling layer. The field u_0^p that the network reconstructs from the coarse-grained DNS snapshot $C(u_0)$ is subsequently advanced in time to compute the loss function (3.7).

training loop, our results reported here are mostly for networks trained with an unroll time of $T_u = 0.22T_e$ (equivalent to 0.25 non-dimensional time units). We suspect the weak influence of T_u to be due to the fact that the neural network only sees a single coarse-grained field, in contrast to 4DVar that optimises to match a specific coarse-grained time series.

The overall architecture of the neural networks depicted in figure 2 is conceptually similar to that applied by P25 in two-dimensional Kolmogorov flow. The networks are purely convolutional and feature a residual network structure (‘ResNet’, He *et al.* 2016): the ResNet ‘layer’ learns a correction to the input, a , which is then added to the input $a \rightarrow \mathcal{F}(a) + a$ (\mathcal{F} is the learned operation). Each ‘Upsampling + residual block’ operation (see figure 2) doubles the size of the input in all three spatial directions. Hence, for a coarse-graining factor $M = 2^n$, we include n individual residual layers to eventually reconstruct a field at the target resolution. All convolutional layers are fully three dimensional and feature 32 filters with either GeLU (Hendrycks & Gimpel 2016) or linear (no) activation functions. The kernel size is fixed at $4 \times 4 \times 4$ with periodic padding applied along all snapshot boundaries. Additional tests performed with a kernel size of $5 \times 5 \times 5$ (which is symmetric in its upstream/downstream coverage around a target point in the velocity field) resulted in almost identical reconstruction errors. The Leray projection discussed above forms the last layer in the model and ensures that the obtained high-resolution field is solenoidal. The entire training process is implemented using the Keras library with a JAX backend (Chollet *et al.* 2015) to allow for straightforward inclusion of the JAX-CFD-based DNS solver and the ability to perform end-to-end differentiation in a single line of code. Note that an extensive variation and optimisation of the model architecture is beyond the scope of this study. As such there will likely be some room for further improvement in the predictive capability of the model.

Optimisation with 4DVar often leads to estimates for the high-resolution state that feature high-wavenumber noise. Early tests of our SRdyn approach revealed similar behaviour, and we found that pre-training the network for 50 epochs to reproduce an interpolated high-resolution field to be highly beneficial to the final model performance




Figure 3. Reconstruction error ϵ_u (see (4.1)) under time advancement for coarsening factors $M = 4$, $M = 8$ and $M = 16$. Reconstructions shown are tri-cubic interpolation (blue), SRdyn (black), 4DVar (orange, solid/dashed/dotted lines for $T_{DA}/T_e = 0.4, 0.9, 1.3$, respectively). The grey (SRdyn) and orange (4DVar, $T_{DA} = 0.9T_e$) shaded regions indicate the ensemble mean \pm one standard deviation over an extended set of trajectories with different, statistically independent initial conditions, while the solid lines are the representative initial condition discussed in much of § 4.

in this context. To this end, the network was first trained on a loss of the form

$$\mathcal{L}_{IC} = \frac{1}{N_S} \sum_{j=1}^{N_S} \|\mathcal{I} \circ \mathcal{C}(\mathbf{u}_j) - \mathcal{N}_\theta \circ \mathcal{C}(\mathbf{u}_j)\|^2, \tag{3.8}$$

where \mathcal{I} represents tri-cubic spline interpolation. In a second step, network training was continued using the SRdyn loss function (3.6), with the pre-trained weights serving as initialisation of the SRdyn network.

4. Results

4.1. Super-resolution with a dynamic loss

We first analyse the predictive performance of the three-dimensional SRdyn technique and compare it with both simple polynomial interpolation and standard 4DVar. Coarse-grained fields, $\mathcal{C}(\mathbf{u}_0)$, from the previously unseen ‘test’ dataset are given to the different state estimation techniques and their reconstructions \mathbf{u}_0^p are compared with the ground truth state. The snapshot reconstructions are obtained as $\mathbf{u}_0^p \equiv \mathcal{N}_\theta \circ \mathcal{C}(\mathbf{u}_0)$ for SRdyn and $\mathbf{u}_0^p \equiv \mathbf{v}_0$ for 4DVar. An initial guess for the 4DVar algorithm is generated by interpolating the coarse-grained field onto the high-resolution grid using tri-cubic interpolation, $\mathcal{I} \circ \mathcal{C}(\mathbf{u}_0)$.

4.1.1. State reconstruction and performance under time marching

To quantify the deviation between predicted and observed trajectories, we first compute the relative reconstruction error

$$\epsilon_u(t) = \frac{\|\boldsymbol{\varphi}_t(\mathbf{u}_0) - \boldsymbol{\varphi}_t(\mathbf{u}_0^p)\|}{\|\boldsymbol{\varphi}_t(\mathbf{u}_0)\|} \tag{4.1}$$

for the various state estimation schemes (4DVar, SRdyn and simple polynomial interpolation). The temporal evolution of (4.1) is reported for the different state estimation schemes in figure 3. The shaded regions in this figure correspond to \pm one standard deviation around the mean for (i) the full test dataset in SRdyn (grey), and (ii) 16 independent initial conditions for 4DVar (orange). The solid lines are for a single initial condition – which is representative of the test set as a whole – that we examine in detail below.




Figure 4. Comparison of the reconstructed velocity fields at $t = 0$ with the high-resolution DNS snapshot that is to be reproduced, alongside the respective coarse-grained fields. Shown are slices of v/u_{rms} at $z = 0$ for the different reconstruction procedures and coarsening factors $M \in \{4, 8, 16\}$, with colours ranging from blue (dark) to yellow (bright) in the interval $[-3, 3]$. For the 4DVar runs, only the best-performing case (in terms of the reconstruction error ϵ_u , cf. (4.1)) is shown for each coarsening factor. Red boxes highlight individual high-speed regions in the ground truth state and its reconstructions.

Consistent with the findings of P25 in two-dimensional Kolmogorov flow, SRdyn provides the most accurate prediction of the initial state and continues to outperform 4DVar for the initial 0.2 to 0.4 large-eddy times, depending on the degree of coarse-graining. With time, the trajectories predicted by classical 4DVar approach the ground truth trajectory and finally reach the lowest reconstruction error around the end of their respective assimilation windows. At these later times, 4DVar strongly benefits from its optimisation, which incorporates observations of the specific ground truth trajectory to which the neural network has no access: the neural network instead generates the most plausible high-resolution state from a single coarse-grained field based on its training experience with a large number of coarse-grained evolutions. In this regard, it is remarkable that the time-marched SRdyn state leads to a comparably good or even better reconstruction of the target trajectory than some of the 4DVar runs, at least for $M = 4$ and $M = 16$. As expected, on a time scale clearly longer than T_{DA} , the reconstruction error of all state estimation techniques rises rapidly and no reliable predictions are possible over these time horizons. Unsurprisingly, tri-cubic interpolation consistently leads to the highest reconstruction errors; it is comparable in performance to 4DVar only at the harshest coarse-graining and then only at $t = 0$.

4.1.2. Instantaneous velocity fields

A visual comparison of reconstructed velocity fields for the example initial condition of figure 3 (solid lines in the figure) is reported in figure 4 for each of the three coarsening factors $M \in \{4, 8, 16\}$. All three techniques appear to do well, at least qualitatively, at the lowest coarse-graining factor $M = 4$, which is unsurprising given the favourable comparison to the synchronisation scale, $\Delta x_c \approx 0.35l_C$. When the coarse-graining factor is

raised to $M = 8 (\Delta x_c \approx 0.69l_C)$ and $M = 16 (\Delta x_c \approx 1.38l_C)$, resolution of typical vortical structures is lost (see figure 1). As a result, tri-cubic interpolation of the under-resolved velocity fields fails to correctly identify the shape and physical location of the intense velocity fluctuations – see the red boxes in figure 4 highlighting particular structures. The learned interpolation operator \mathcal{N}_θ reduces the exaggerated smoothing effect of simple interpolation and retains – in contrast to both 4DVar and polynomial interpolation – the ability to correctly reconstruct the location of the individual intense velocity patterns (examples highlighted in red boxes), even beyond the synchronisation limit l_C . The 4DVar reconstructions are increasingly dominated by high-wavenumber oscillations and feature clearly overestimated peak values at the larger coarse-graining values.

Corresponding probability density functions (PDFs) of the velocity, vorticity and dissipation rate are shown in figure 5 (upper rows in each panel (a, b, c) correspond to the snapshots of figure 4, lower rows are computed at $t = T_e$). For the mild coarse-graining $M = 4$, all three techniques reproduce the PDF of the velocity fluctuations u' at $t = 0$ almost perfectly, with slight deviations in the tails of the 4DVar-based reconstruction. On the other hand, 4DVar overestimates the tails of the PDFs for both vorticity fluctuations ω' and the local dissipation rate ε , while underestimating the more probable events. Classical polynomial interpolation and SRdyn provide a much better reconstruction of the ground truth PDFs, with only the most extreme values being under-represented in the reconstructed field.

Consistent with the observations in figure 4, 4DVar’s tendency to overestimate the PDF of the velocity fluctuations increases with further coarse-graining, while SRdyn (and also the simple polynomial interpolation) tends to smooth out the extreme values in the field more strongly. The same trend is seen in the PDFs of ω' and ε , but the deviations from the target PDF are more pronounced. The distributions of the 4DVar-reconstructed fields are much wider than that of the ground truth state – the assimilated fields overestimate the occurrence of extreme gradients. On the other hand, the low-amplitude values are over-represented in the fields generated via simple polynomial interpolation. The super-resolution approach SRdyn consistently provides the most faithful reconstruction of the PDFs, although its deviations are still significant for $M = 16$. Given the poor performance of simple tri-cubic interpolation across the various metrics analysed so far, we do not consider this method further beyond this point.

Under time advancement, all PDFs evolve towards the ground truth as the corresponding trajectories rapidly collapse onto the turbulent attractor and the initial unphysical oscillations of 4DVar decay. After one large-eddy time (lower rows of each panel in figure 5), all PDFs essentially collapse onto the ground truth for $M = 4$ and $M = 8$. For $M = 16$, deviations especially in the PDF tails remain visible, though they are much weaker than for the initial reconstructed fields u_0^p .

In order to quantify how well the different scales of the velocity field are reconstructed by the different state estimation techniques, we introduce the normalised co-spectrum of u and u^p (Li *et al.* 2020):

$$\widehat{\rho}_{uu^p}(k, t) = \frac{\int_{|k|=k} d\mathbf{k} \hat{u}(\mathbf{k}, t) \cdot \hat{u}^p(\mathbf{k}, t)^*}{2\sqrt{E_u(k, t)E_{u^p}(k, t)}}. \tag{4.2}$$

Here E_u and E_{u^p} are the energy spectra for a ground truth state and its reconstruction, respectively. Hence, $\widehat{\rho}_{uu^p}$ measures the correlation between a ground truth Fourier mode \hat{u} and its reconstructed counterpart \hat{u}^p for each wavenumber k , with a value close to unity (zero) indicating a very good (poor) reconstruction of the respective Fourier mode.

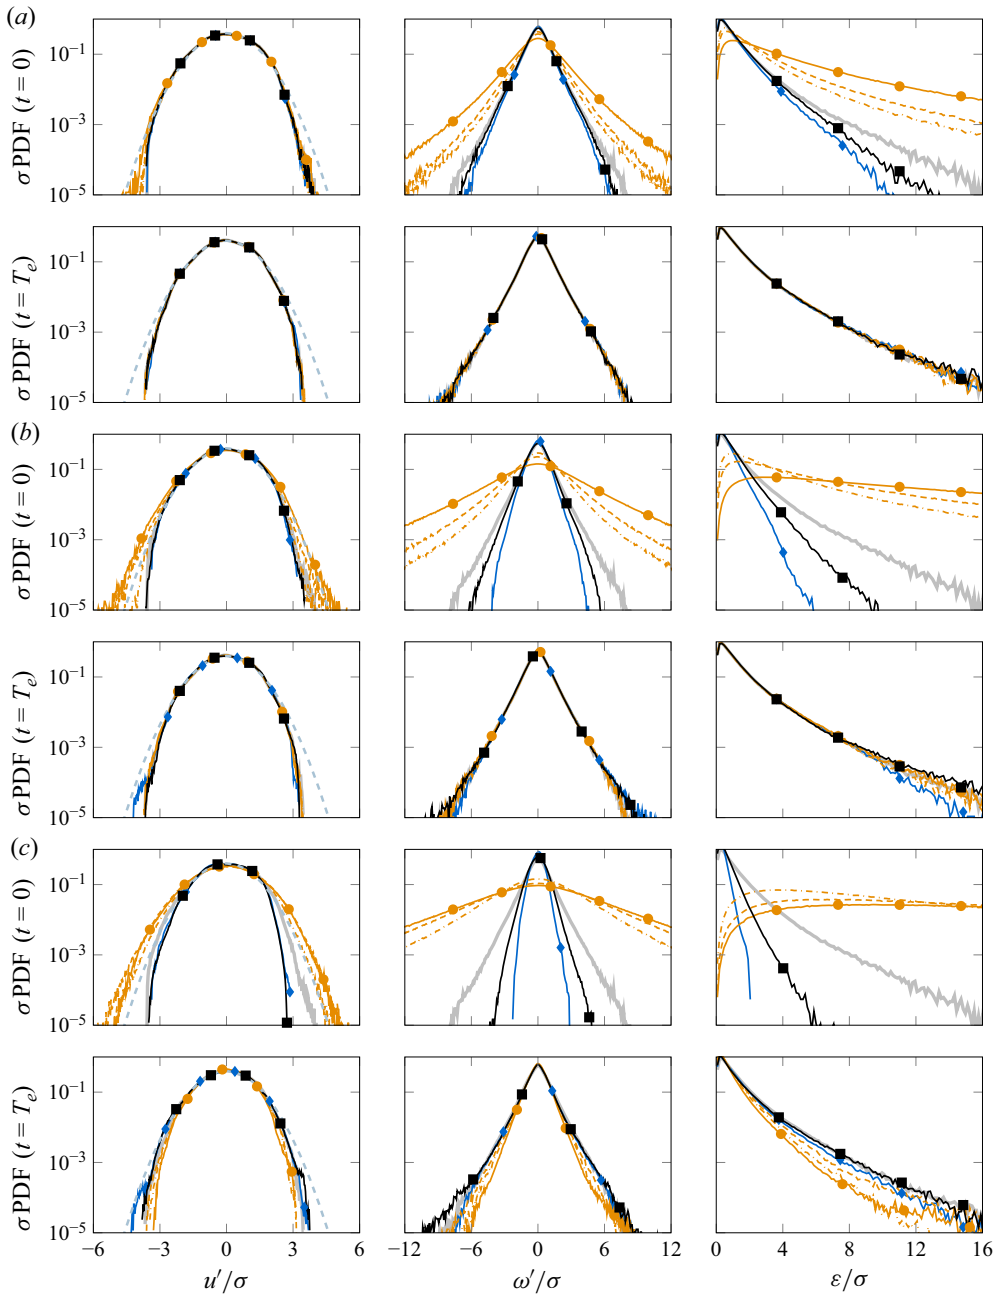


Figure 5. Standard PDFs of (left) the velocity fluctuations u' ($\sigma = u_{rms}$), (middle) the vorticity fluctuations ω' ($\sigma = \omega_{rms}$) and (right) the local dissipation rate ϵ ($\sigma = \bar{\epsilon}$) for the reconstructed fields at $t=0$ (rows one, three and five) and $t=T_e$ (rows two, four and six): (a) $M=4$, (b) $M=8$, (c) $M=16$. Variables shown are the DNS ground truth (thick dark grey), tri-cubic interpolation (blue), SRdyn (black), 4DVar (orange, line styles as before) and the standard normal distribution (light blue, dashed, shown for u' only).




Figure 6. Normalised co-spectrum $\hat{\rho}_{uu^p}(k, t)$ between the ground truth \mathbf{u} and the reconstructed field \mathbf{u}^p for the SRdyn (black) and 4DVar runs (orange, line styles as before) as a function of the wavenumber at $t/T_e \in \{0, 0.5, 1\}$ for coarsening factors (a) $M = 4$, (b) $M = 8$ and (c) $M = 16$. Vertical lines indicate the maximum forcing wavenumber $k_f = 3$ (solid, grey), the Nyquist cutoff wavenumber k_N of the respective coarsening (dashed, purple) and the critical wavenumber $k_C = 0.2\eta_K^{-1}$ (dash-dotted, blue), respectively.

In figure 6, $\hat{\rho}_{uu^p}$ is evaluated for the different coarse-graining factors at times $t/T_e \in \{0, 0.5, 1\}$ along the trajectories. The Nyquist cutoff wavenumbers of the coarse-graining operator are larger than the synchronisation value at $M = 4$ and $M = 8$ ($k_N \approx 2.9k_C$ and $k_N \approx 1.4k_C$, respectively), but smaller at $M = 16$ ($k_N \approx 0.7k_C$). In the SRdyn-generated reconstruction, all scales up to the Nyquist wavenumber are more or less perfectly reproduced for the different coarsening factors, followed by a range of partially recovered modes $k > k_N$ whose number reduces with increasing M . On the other hand, 4DVar has difficulty reconstructing the large scales of the initial velocity field, with markedly reduced values of $\hat{\rho}_{uu^p}$ appearing before the Nyquist cutoff. For the strongest coarse-graining $M = 16$, standard 4DVar even struggles to accurately reproduce the largest scales into which energy is injected by the external forcing.

As the reconstructed fields are unrolled for $M = 4$ (figure 6a), both techniques provide accurate reproductions of all dynamically relevant scales, with SRdyn even outperforming some of the 4DVar runs up to $t = T_e$. For $M = 8$ (figure 6b), SRdyn is still able to faithfully reproduce the large-scale content of the velocity field over a large-eddy time, but the high-wavenumber modes deviate from those in the ground truth trajectory as time increases. This is not unexpected since the smallest scales are known to decorrelate much faster than their large-scale counterparts (Boffetta & Musacchio 2017). The 4DVar runs, on the other hand, are optimised to match the coarse-grained observations as closely as possible and

provide better approximations for a wider range of scales when $t \geq 0.5T_e$. Finally, with k_N falling well below k_C at $M = 16$ (figure 6c), even 4DVar is not capable of reproducing much more than the largest flow scales in a satisfactory way. Notably, the high-resolution fields obtained by SRdyn again outperform all of their 4DVar-generated counterparts in the early stages of the time evolution and continue performing equally well as the 4DVar run with the shortest assimilation window, i.e. $T_{DA} \approx 0.4T_e$.

Our observations on the performance of 4DVar, in particular the decreasing quality with increasing coarse-graining, is in line with what is known about assimilation in the context of turbulence synchronisation (see discussion in § 1) and is in qualitative agreement with the results of Li *et al.* (2020) who performed a similar 4DVar-based study in a triply periodic flow with monochromatic ‘Kolmogorov’ forcing (which drives a large-scale mean flow). A one-to-one comparison with the latter study is not straightforward due to the change in forcing profile and the fact that their 4DVar algorithm operates in spectral space, where the leading Fourier modes of the ground truth at wavenumbers $0 \leq k \leq k_O$ act as observations for the optimisation. For comparison, we performed some 4DVar runs for the Kolmogorov forcing used by Li *et al.* (2020) at identical parameters – the key difference is that the optimisation is performed in physical space and exclusively with coarse-grained observational data. These calculations are included in Appendix B. In line with previous observations in other data assimilation techniques (Di Leoni, Mazzino & Biferale 2020), the comparison reveals that, for a Nyquist cutoff k_N matching the spectral cutoff wavenumber k_O , 4DVar based on coarse-grained physical observations leads to weaker reconstructions than the purely spectral procedure of Li *et al.* (2020).

4.1.3. Reproduction of vortical structures

Snapshots of the instantaneous vorticity field derived from the reconstructed state, $\omega^P(\mathbf{x}, t) := \nabla \times \boldsymbol{\varphi}_t(\mathbf{u}_0^P)$, are shown in figure 7 for the intermediate coarse-graining level $M = 8$ alongside the corresponding ground truth states. Given that both neural network training and data assimilation have been performed based on velocity field snapshots, reconstructing gradient-based fields like ω^P is an additional challenge. The SRdyn-based reconstruction of the initial vorticity field $\omega_0^P := \omega^P(\mathbf{x}, 0)$ reveals a qualitatively similar pattern of small-scale structures to the target DNS snapshot, albeit contaminated with some background noise. The vorticity field reconstructed by standard 4DVar, on the other hand, deviates significantly from the target state in terms of both the spatial organisation and amplitude of the intense vorticity regions, consistent with the overestimated tails of the corresponding PDF in figure 5.

When advanced in time, the strong initial deviations between the 4DVar-reconstructed field and the DNS reference state decay rapidly, as does the weaker background noise in the SRdyn-based reconstruction. At $t = 0.25T_e$, the deviations of the SRdyn-based reconstruction and the ground truth are minimal (see the highlighted regions in figure 7). While 4DVar also provides a reasonable reconstruction of the vorticity field, the deviations from the ground truth are more clearly discernible. This situation changes at $t = 0.5T_e$, where 4DVar now provides a slightly more accurate reconstruction of the ground truth, though the SRdyn-based reconstructions still perform well considering that they have not seen observational data of the specific trajectory. For $t > 0.5T_e$, deviations between the SRdyn-generated reconstructions and the ground truth snapshots grow stronger – at $t = T_e$, only the large-scale organisation of the vorticity fields are still comparable. A similar deviation from the ground truth trajectory occurs for the 4DVar-based reconstruction once the trajectory leaves the assimilation window (see snapshots at $t = 2T_e$ in figure 7).




Figure 7. Comparison of the unrolled vorticity trajectory initialised with the reconstructed vorticity field ω_0^p for SRdyn and 4DVar ($T_{DA} \approx 1.3T_e$) and a coarsening factor $M = 8$. Shown are slices of ω_y/ω_{rms} at $z = 0$, with colouring as in figure 1. Red boxes highlight individual structures in the ground truth state and their reconstructions.

A complementary view is provided in figure 8, where the three-dimensional organisation of vortical structures and regions of intense strain is visualised using isosurfaces of the second invariant of the velocity gradient tensor Q (Hunt, Wray & Moin 1988). In the figure, red ($Q > 2.5\omega_{rms}^2$) and blue ($Q < -2.5\omega_{rms}^2$) isosurfaces indicate regions of strong vortical activity and large strain, respectively. In agreement with our observations of vorticity slices in figure 7, SRdyn provides a reasonable reconstruction of the three-dimensional vortical structures in the initial field, though the field is less smooth than the ground truth due to the presence of low-amplitude background noise. For 4DVar, the range of values attained by Q is seen to be greatly overestimated so that no individual features are visible with isosurfaces at $Q = \pm 2.5\omega_{rms}^2$. The initial low-amplitude reconstruction errors in the SRdyn-generated state rapidly diffuse under time advancement and, for $t = 0.25T_e$ and $t = 0.5T_e$, most vortices and intense-strain regions in the DNS target field are well captured by the neural network prediction (see the regions highlighted with green boxes in figure 8). For longer unroll times, the overall spatial organisation of the vortices remains comparable to the ground truth DNS, but the small-scale features in the SRdyn-based reconstruction start to deviate. The high errors of the 4DVar reconstruction also dissipate rapidly as the initial field is unrolled in time, and by $t = 0.25T_e$ the Q isosurfaces largely match those in the reference trajectory. However, it takes until $t = 0.5T_e$ for the 4DVar-based prediction to outperform SRdyn. Towards the end of the assimilation window at $T_{DA} \approx 1.3T_e$, 4DVar reaches its best predictive performance (compare the minimum of ϵ_u in figure 3) and deviations from the ground truth are hardly visible.

Figure 9 compares the corresponding reconstructions of the joint probability density function (JPDF) for the second (Q) and third (R) invariant of the velocity gradient tensor with those of the ground truth DNS trajectory. Consistent with figure 8, it is seen that the 4DVar-based reconstruction of the JPDF initially clearly deviates from the characteristic

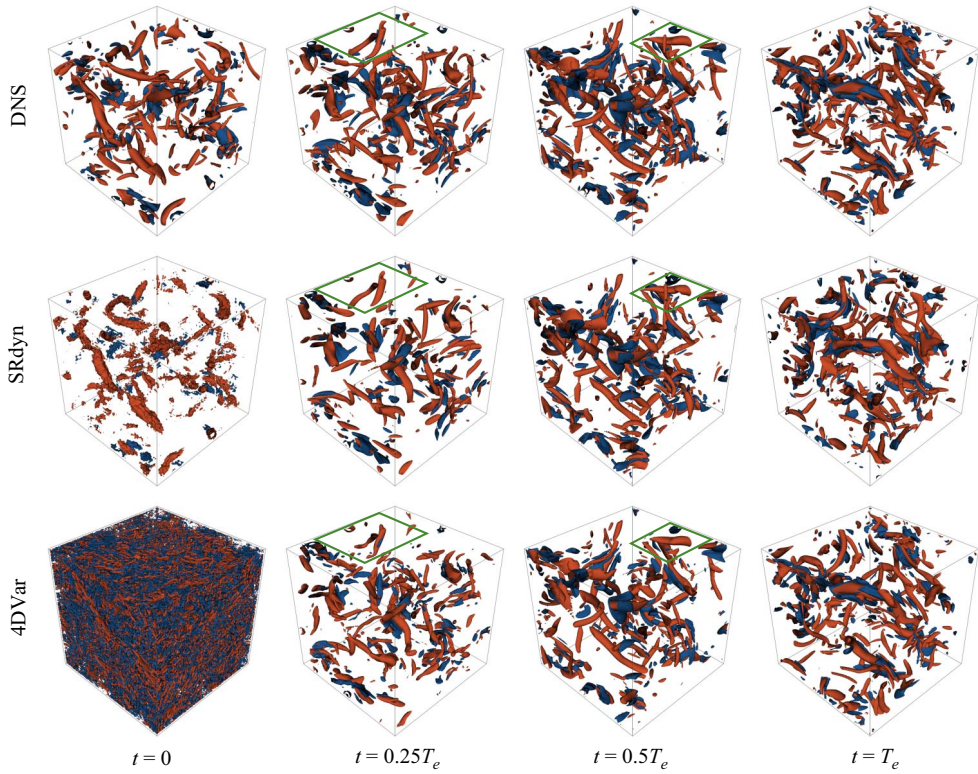


Figure 8. Evolution of vortical structures (red) and regions of large squared strain (blue) from state estimation at $M=8$, visualised using isosurfaces of the second invariant of the velocity gradient tensor at thresholds $Q/\omega_{rms}^2 = -2.5$ (blue) and $Q/\omega_{rms}^2 = +2.5$ (red). The 4DVar results are obtained with an optimisation over a window of length $T_{DA} = 1.3T_e$. Green rectangles highlight individual vortical structures and their reconstructions.

teardrop-like shape in the DNS that indicates a dominance of vortex stretching ($Q > 0$, $R < 0$) and strain-rate self-amplification events ($Q < 0$, $R > 0$). In addition, the range of values attained by Q and R is strongly over-estimated. On the other hand, SRdyn generates reasonable reproductions of the ground truth state statistics from the beginning, missing only some of the most extreme values. The 4DVar-based reconstruction successively improves with time and, for $t \geq 0.5T_e$, SRdyn and 4DVar both reveal similar distributions as in the DNS ground truth.

4.2. 4DVarSR– a hybrid approach

The results above show that the SRdyn approach – which does not rely on the availability of a library of high-resolution data – can outperform 4DVar for state estimation in three-dimensional turbulence at $t = 0$ (as measured by an L_2 norm to the ground truth), in line with the results reported in P25 for two-dimensional Kolmogorov flow.

Similar to the two-dimensional work, this performance does not continue to hold under time marching: after an interval where both methods provide comparably good reconstructions of the target trajectory, standard 4DVar eventually wins thanks to its access to observations of the dynamics along the specific trajectory. A well-known challenge in the nonlinear optimisation framework of 4DVar is the non-uniqueness of the

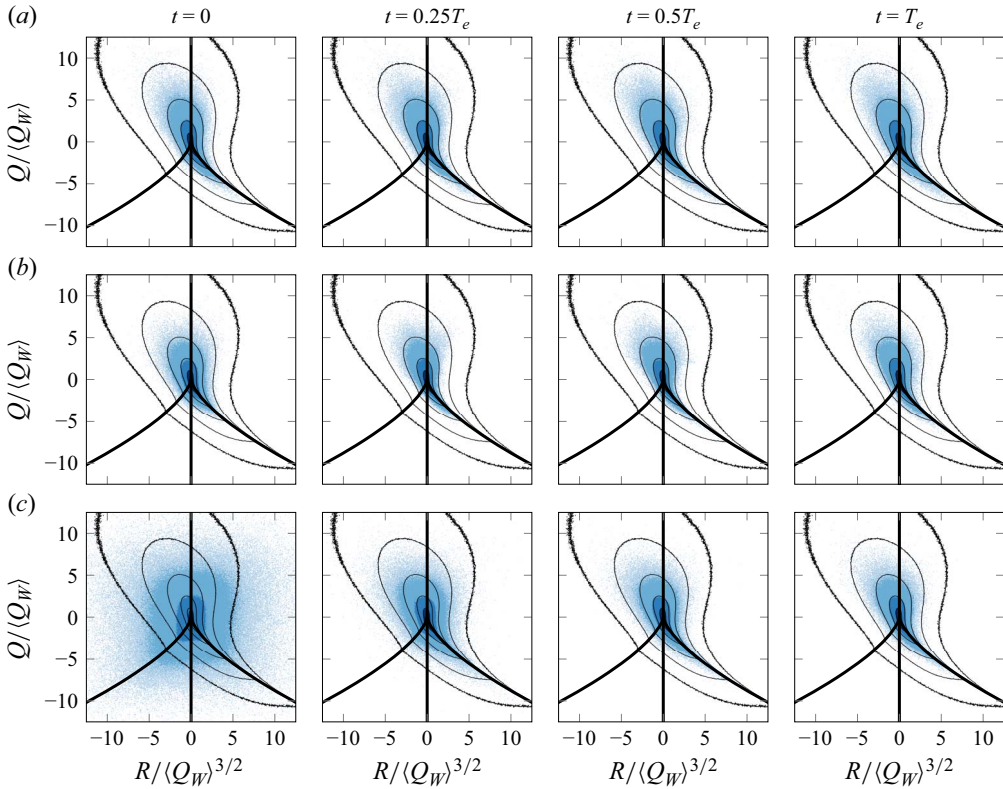


Figure 9. Evolution of the standard JPDF of the second (Q) and third (R) invariant of the velocity gradient tensor, normalised with the spatio-temporal mean of $Q_w = (\omega_i \omega_i)/4$, for (a) the ground truth, (b) SRdyn and (c) standard 4DVar. Contour levels are spaced logarithmically with a minimum value of 10^{-5} and colours ranging from light to dark blue. Thin black contour lines represent ensemble-averaged long-time statistics over a set of 100 independent DNS trajectories, thick black curves are $R = 0$ and $Q^3 + 27R^2/4 = 0$.

target high-resolution field – several trajectories starting from very different states \mathbf{u}_0^p may reconstruct the coarse observations at late times with similar accuracy (Zaki 2025).

One approach to improve the poor reconstruction of early time states in 4DVar is to emphasise reconstruction of early observations by adding time-dependent weights to the loss function (Wang *et al.* 2019). Here, we instead aim at improving the quality of the prediction by initialising the 4DVar optimisation algorithm with a more accurate, neural network-generated reconstruction of \mathbf{u}_0 as the initial guess, all without requiring any knowledge of high-resolution fields. Our approach (hereafter termed ‘4DVarSR’) shares some similarities with the method of Frerix *et al.* (2021), who used a neural network trained on high-resolution data in two-dimensional Kolmogorov flows to initialise 4DVar. That approach has an advantage in that the networks require a much shorter training time (due to lack of time marching) but, similar to the motivation behind the SRdyn algorithm discussion in § 3, is perhaps suboptimal when working directly with experimental data where generation of a full high-resolution dataset may be inappropriate.

The positive influence of the SRdyn-based initialisation on the 4DVar reconstruction of \mathbf{u}_0 is demonstrated in figure 10, where the 4DVarSR field \mathbf{u}_0^p is shown alongside the results obtained with SRdyn and 4DVar at the more challenging coarse-graining levels




Figure 10. Comparison of the 4DVarSR-reconstructed velocity fields at $t = 0$ with the high-resolution DNS snapshot that is to be reproduced, the respective coarse-grained fields and those obtained with straight SRdyn and 4DVar. Shown are slices of v/u_{rms} at $z = 0$ for the different reconstruction procedures and coarsening factors $M \in \{8, 16\}$, with colouring as in figure 4. For the 4DVar and 4DVarSR runs, only the best-performing case is shown for each value of M . Position and size of the red boxes that highlight selected high-speed regions are identical to those in figure 4.

$M = 8$ and $M = 16$. The 4DVarSR-generated prediction features much less contamination with erroneous small-scale features than in standard 4DVar, while the predicted values fall in a similar range as those in the ground truth. The hybrid 4DVarSR approach is also able to correctly locate the high-speed regions marked by red boxes for $M = 16$, in contrast to standard 4DVar that predicts a state in which the most intense velocity patches are shifted with respect to the DNS ground truth. At the less-demanding value of $M = 8$, a comparison of the 4DVarSR- with the SRdyn-based reconstruction reveals only minor deviations. This is not unexpected since the SRdyn is used to initialise the 4DVarSR procedure. However, at the strongest coarse-graining $M = 16$, the 4DVarSR high-resolution snapshot \mathbf{u}_0^p differs more strongly from SRdyn. Some of the erroneous small-scale features characteristic of 4DVar have emerged during the optimisation, but at a much weaker amplitude so that they do not dominate the reconstructed field.

Figure 11 presents standardised PDFs of velocity and vorticity fluctuations as well as fluctuations of the local dissipation rate ε for $M = 8$ shown in subfigure (a) and $M = 16$ in subfigure (b). In both subfigures, the upper rows correspond to data at $t = 0$ and the bottom rows to data at $t = T_e$. In support of the previous observations, the 4DVarSR prediction of the PDFs at $t = 0$ represents a significant improvement on those obtained with ‘standard’ 4DVar and SRdyn (latter not shown here, see figure 5). The results are particularly impressive for the two gradient-based quantities ω' and ε , for which standard 4DVar greatly over-predicts the tails of the distribution and SRdyn misses the most extreme values. The 4DVarSR PDFs almost perfectly reproduce the DNS ground truth data at the intermediate coarse-graining level $M = 8$, and clearly reduce the over-prediction of the PDF tails compared with the standard 4DVar at the more demanding $M = 16$. Unrolling the reconstructed fields over one large-eddy time leads to an almost perfect match between 4DVarSR and the PDFs of the ground truth even at $M = 16$, while standard 4DVar still underestimates the probability of intense vorticity and dissipation events, as can be seen from the data shown in rows two and four of figure 11.

In figure 12 the evolution of the pointwise relative reconstruction error ϵ_u is compared for 4DVar and 4DVarSR for the example trajectory. Initialising the data assimilation procedure with the neural network prediction reduces the initial error $\epsilon_u(0)$ by more




Figure 11. Standard PDFs of (left) the velocity fluctuations u' ($\sigma = u_{rms}$), (middle) the vorticity fluctuations ω' ($\sigma = \omega_{rms}$) and (right) the local dissipation rate ε ($\sigma = \bar{\varepsilon}$) for the reconstructed fields at $t=0$ (rows one and three) and $t=T_e$ (rows two and four): (a) $M=8$, (b) $M=16$. Variables shown are the DNS ground truth (thick dark grey), 4DVar (orange, line styles as before), the best-performing 4DVarSR run (red) and the standard normal distribution (light blue, dashed, shown for u' only).




Figure 12. Comparison of the model predictions under time advancement in terms of the reconstruction error ϵ_u for coarsening factors $M=4$, $M=8$ and $M=16$. Reconstructions shown are 4DVar (orange, line styles as before) and the best-performing 4DVarSR run (red).

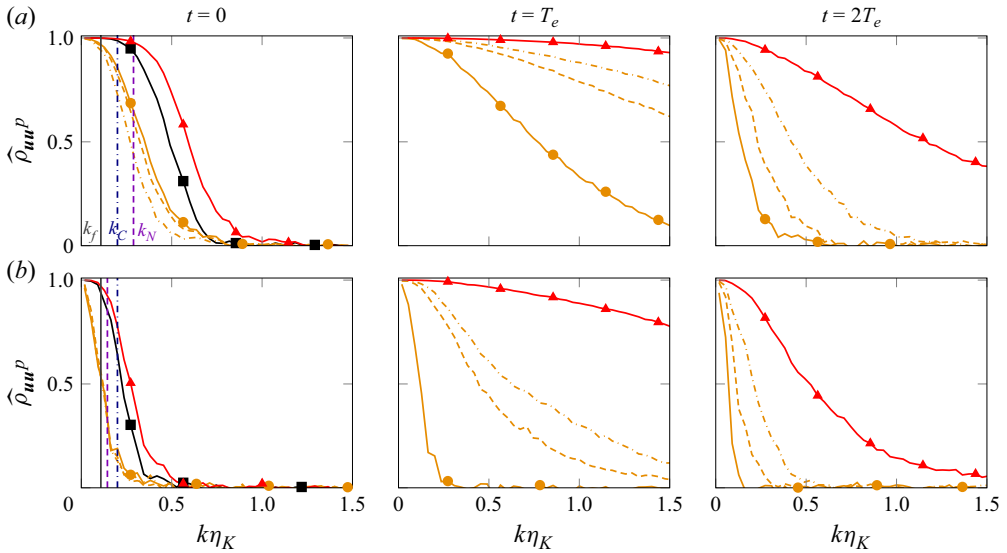


Figure 13. Normalised co-spectrum $\hat{\rho}_{uu^p}(k, t)$ between the ground truth \mathbf{u} and the reconstructed field \mathbf{u}^p at $t/T_e \in \{0, 1, 2\}$ for coarsening factors (a) $M = 8$ and (b) $M = 16$. Shown here are 4DVar (orange, line styles as before) and the best-performing 4DVarSR run (red). At $t = 0$, the SRdyn-based reconstruction (black) is shown for comparison as well. Vertical lines are as in figure 6.

than a factor two compared with standard 4DVar. Beyond $t = 0$, the initialisation with the SRdyn-generated field has a lasting influence on the quality of the prediction over the entire observation interval, allowing 4DVarSR to outperform the classical 4DVar predictions for all coarse-graining levels and assimilation window lengths. Moreover, for both $M = 4$ and $M = 8$, the error growth rate after passing the overall minimum is noticeably reduced compared with the 4DVar runs – reliable predictions are possible over longer time intervals.

The Fourier correlation coefficient $\hat{\rho}_{uu^p}(k, t)$ reported in figure 13 shows that the significantly lower reconstruction errors in 4DVarSR are associated with a more accurate reproduction of a wider range of wavenumbers in the initial field \mathbf{u}_0 than both SRdyn and ‘standard’ 4DVar. The 4DVarSR method has further improved the SRdyn prediction of \mathbf{u}_0 , providing more faithful reconstructions of the modes with $k > k_N$. As the 4DVarSR fields are unrolled in time, the resulting trajectory follows the ground truth more robustly than 4DVar: even at $t = T_e$, we observe mode-wise correlations of more than 75 % for all resolved wavenumbers. Later at $t = 2T_e$, which is well beyond the end of the assimilation window $T_{DA} \approx 1.3T_e$, the 4DVarSR reconstruction of the low-wavenumber content is still fairly good at both coarse-graining levels. For $M = 8$, this holds even for most of the high-wavenumber modes, and snapshots of both the velocity and vorticity fields are still almost indistinguishable from their ground truth counterparts (not shown).

For the strongest coarse-graining, figure 14 shows selected snapshots of the 4DVar- and 4DVarSR-generated vorticity ω^p , alongside the corresponding states of the ground truth DNS. The initial field reconstructed by means of standard 4DVar shares very little similarity with the ground truth, with the range of reconstructed values far exceeding that of the target field. The range of values attained in the 4DVarSR-reconstructed field is much closer to that of the DNS reference state, but still features a pronounced unphysical high-wavenumber noise, and only the large-scale pattern roughly matches the overall flow organisation in the target field. For standard 4DVar, the vorticity amplitudes decay quickly




Figure 14. Comparison of the unrolled vorticity trajectory for ‘standard’ 4DVar and 4DVarSR (both $T_{DA} \approx 1.3T_e$) and a coarsening factor $M = 16$. Shown are slices of ω_y/ω_{rms} at $z = 0$, with colouring as in figure 1. Red boxes highlight individual structures in the ground truth state and their reconstructions.

with time and reach a range comparable to the ground truth state by $t = 0.25T_e$, though the spatial distribution of ω_y^p remains fairly different from the corresponding ground truth state until $t = T_e$. Beyond the end of the assimilation window, the 4DVar-based reconstruction deviates strongly from the target field as the assimilated trajectory diverges from its ground truth counterpart. For 4DVarSR, on the other hand, the initial high-wavenumber noise has diffused away by $t = 0.25T_e$ and, from $t = 0.5T_e$ on, the vorticity field is visually almost indistinguishable from the ground truth up to a large-eddy time.

Figure 15 shows the spatial distribution of the local reconstruction error $\omega_y^p(\mathbf{x}, t) - \omega_y(\mathbf{x}, t)$ corresponding to the snapshots in figure 14. While standard 4DVar features errors across the entire domain for all snapshots (except for reasonable agreement at $t = T_e$), the overall error amplitude for 4DVarSR is much lower. Deviations from the ground truth occur primarily in regions of intense vorticity. A comparison with the ‘full’ ground truth snapshots ω_y indicates that 4DVarSR reproduces the spatial organisation of the vorticity field as a whole and that the error seems to originate in a mild under- or over-prediction of the amplitude in the patches of intense vorticity. Beyond the end of the assimilation window (see snapshots at $t = 2T_e$), we find that the reconstruction error for 4DVarSR does not grow uniformly in space. Instead, individual flow features such as the highlighted intense vorticity patches are no longer captured in the reconstruction; for other flow features, the main error contribution seems to be an incorrectly predicted amplitude.

Figure 16 presents a comparison of SRdyn, 4DVar and 4DVarSR with respect to their ability to faithfully reproduce the turbulent energy spectra of the target snapshots, with results pertaining to $M = 8$ and $M = 16$ shown in subfigures (a) and (b), respectively. Results for $t = 0$ are shown in the leftmost panels of both subfigures. As can be seen from the data shown in there, SRdyn under-predicts the energy for $k_N < k < 1/\eta_K$, while 4DVar over-predicts the energy levels for these wavenumbers. The 4DVarSR-based spectra fall in between and provide the best reconstruction of the ground truth state, mildly




Figure 15. Comparison of the local reconstruction error $(\omega_y^p - \omega_y)/\omega_{rms}$ at $z = 0$ along the unrolled vorticity trajectories for 4DVar and 4DVarSR (both $T_{DA} \approx 1.3T_e$) and a coarsening factor $M = 16$. For the sake of comparison, ‘full’ snapshots ω_y/ω_{rms} from the ground truth DNS trajectory are included in the top row. Colours range from blue (dark) to yellow (light) in the interval $[-4, 4]$. Red boxes highlight individual structures in the ground truth state and the respective local deviations in the reconstructed states.




Figure 16. Instantaneous normalised energy spectra as a function of the wavenumber for snapshots at $t/T_e \in \{0, 1, 1.5\}$ and coarsening factors (a) $M = 8$ and (b) $M = 16$. Shown here are the DNS ground truth (thick grey), SRdyn (black), standard 4DVar (orange, line styles as before) and the best-performing 4DVarSR run (red). Vertical lines are as in figure 6.




Figure 17. Instantaneous normalised energy fluxes, $\Pi/\bar{\varepsilon}$, as a function of the wavenumber for snapshots at $t/T_e \in \{0, 1, 1.5\}$ and coarsening factors (a) $M = 8$ and (b) $M = 16$. Shown here are the DNS ground truth (thick grey), SRdyn (black), standard 4DVar (orange, line styles as before) and the best-performing 4DVarSR run (red). Vertical lines are as in figure 6 and the dashed horizontal line marks $\Pi = \bar{\varepsilon}$.

under-predicting the energy contributions at wavenumbers just above k_N and slightly over-predicting in the dissipative range. The high- k over-estimation in 4DVarSR and SRdyn remains well below that of 4DVar. This is consistent with our previous observations that 4DVarSR significantly reduces the intense high-wavenumber noise in the initial field that was seen to characterise standard 4DVar at more intense coarse-graining. The middle and right panels of figure 16 correspond to measurements later in time, at $t = T_e$ and $t = 1.5T_e$, respectively. As can be seen from the data shown in these panels, the initial deviations in the mid- to high-wavenumber modes present in all reconstructions at $t = 0$ quickly disappear as the fields are unrolled in time. While the reconstructed spectra coincide with that of the ground truth at later times for $M = 8$, both SRdyn and 4DVar are seen to somewhat deviate from the target spectrum for $M = 16$. Again 4DVarSR shows a better predictive performance and reproduces the ground truth almost perfectly, even for the most severe coarsening.

Figure 17 shows the corresponding instantaneous energy flux across wavenumber k , $\Pi(k, t)$, at three instances in time $t/T_e \in \{0, 1, 1.5\}$ for $M = 8$ (subfigure a) and $M = 16$ (subfigure b). As can be seen from the leftmost panel of subfigure (a), for the initial instant in time at the moderate coarse-graining level $M = 8$, both standard 4DVar with $T_{DA}/T_e \in \{0.9, 1.3\}$ and SRdyn reproduce the general shape of the flux for $k \leq k_N$. In this context, standard 4DVar with $T_{DA}/T_e = 0.9$ and $T_{DA}/T_e = 1.3$ is seen to perform well even at higher wavenumbers. Nonetheless, in comparison with the ground truth, the inter-scale energy flux is under-predicted by all methods in different wavenumber ranges (all wavenumbers in the case of standard 4DVar at $T_{DA} \approx 0.4T_e$). The 4DVarSR-based reconstruction of Π coincides very well with the DNS ground truth up to the Nyquist cutoff wavenumber k_N , while predicting too low energy transfer rates towards higher wavenumbers.

The leftmost panel of figure 17(b) presents results at $t = 0$ for $M = 16$. As can be seen from the data shown in the figure, 4DVarSR (shown in red) is the only technique that is

able to roughly reproduce the characteristic shape of the ground truth flux, coinciding with the reference DNS for all $k < k_N$. The 4DVar and SRdyn results, on the other hand, differ significantly in shape and amplitude from the ground truth. We find that SRdyn strongly under-predicts the energy transfer to higher wavenumbers over the entire wavenumber range, while standard 4DVar greatly over-estimates Π for the lower wavenumbers and, in some cases, even turns negative for the higher wavenumbers, thereby predicting unphysical upscale energy transfer that is absent in the ground truth data.

As for the energy spectra, the reconstructed energy fluxes are seen to approach the ground truth as the initial fields are advanced in time for both coarse-graining levels, as can be seen from the data shown in the middle and right panels of [figure 17](#). For $M = 8$, all methods reproduce the ground truth fluxes rather well. However, for the strongest coarsening $M = 16$ with data shown in subfigure (b), 4DVarSR clearly outperforms the remaining techniques, providing very accurate reproductions of the ground truth flux even beyond the end of its assimilation window, where standard 4DVar and SRdyn retain visible deviations from the ground truth.

4.3. Performance at higher Re_λ

A central weakness of neural networks is that their performance can rapidly deteriorate when exposed to data outside of the training distribution. By identifying ‘invariant’ structures in decaying three-dimensional turbulence via a data-driven, Buckingham-Pi based rescaling of Q and R distributions, Fukami, Goto & Taira (2024) were able to predict where a (super-resolution) network trained at one Re_λ value would be able to make sensible predictions at other Reynolds numbers. Such an assessment of our networks is not possible here due to our fixed Re_λ , though we are able to show how our models can be combined with data assimilation to make extremely accurate predictions at unseen parameter settings, in particular at a Reynolds number larger than that of the original training data.

To do this, a second database of DNS snapshots at $Re_\lambda \approx 112$ (equivalent to $Re = 125$) has been compiled. The resolution has been increased to $N_x \times N_y \times N_z = 256^3$ collocation points in order to maintain a maximum resolved wavenumber $k_m \eta_K = 1.5$. The synchronisation scale $l_C \approx 11.5 \Delta x$ is approximately the same as in the previous sections when measured in terms of the grid spacing, but reduces by a factor two compared with the largest scales. Recalling that all networks were trained to map a coarse-grained snapshot onto a target field of size $N_x \times N_y \times N_z = 128^3$, we proceed as follows: first, a higher-Reynolds-number snapshot is coarsened by a factor M . In a second step, the corresponding network trained at $Re_\lambda \approx 70$ for a coarse-graining factor $M/2$ is used – without further training on the new snapshots – to generate a reconstruction of the ground truth with a size of $N_x \times N_y \times N_z = 128^3$. Finally, the resulting field is interpolated onto the target grid $N_x \times N_y \times N_z = 256^3$ using a tri-cubic interpolation scheme.

Focusing initially on the neural network reconstruction, [figure 18](#) compares the resulting velocity fields \mathbf{u}^p for one representative snapshot with those obtained by classical tri-cubic interpolation for coarse-graining factors $M = 8$ ($\Delta x_c = 0.7l_C$) and $M = 16$ ($\Delta x_c = 1.4l_C$). As for the lower-Reynolds-number cases discussed earlier in this paper (see [figure 4](#)), SRdyn generates a visibly more accurate reconstruction of the target snapshot than standard polynomial interpolation. It suffers much less from blurring and is, even at a coarse-graining level about 40 % above the synchronisation limit, able to correctly locate relevant intense flow structures, though successively under-predicting their peak values (compare the red boxes in [figure 18](#)).

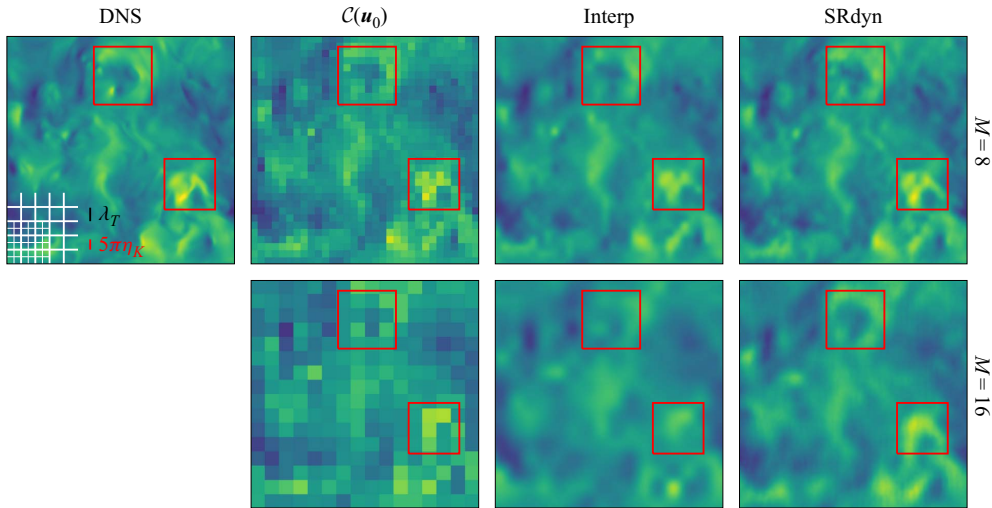


Figure 18. Comparison of the reconstructed velocity fields at $t = 0$ with the high-resolution DNS snapshot at $Re_\lambda \approx 112$ that is to be reproduced, alongside the respective coarse-grained fields. Shown are slices of v/u_{rms} at $z = 0$ for the different reconstruction procedures and coarsening factors $M \in \{8, 16\}$ with respect to the 256^3 grid, with colouring as in figure 4. Red boxes highlight individual high-speed regions in the ground truth state and its reconstructions. White thin and thick grid lines in the lower left corner of the first frame visualise the coarsened grids for coarsening factors $M \in \{8, 16\}$ relative to the DNS grid, in comparison with the Taylor microscale λ_T and the synchronisation limit l_C .




Figure 19. Reconstruction error ϵ_u under time advancement for 4DVar (orange) and 4DVarSR (red) and coarsening factors $M = 8$ and $M = 16$ with respect to the 256^3 grid at $Re_\lambda \approx 112$. Lines indicate ensemble means over four distinct trajectories with different, statistically independent initial conditions for assimilation windows of lengths $T_{DA}/T_e = 0.21$ (solid) and $T_{DA}/T_e = 0.43$ (dashed). Shaded regions indicate the ensemble mean \pm one standard deviation.

Moreover, consistent with our observations in §4.2, the more accurate predictions obtained with the neural networks serve as good initial conditions for data assimilation at the much higher Re_λ value. Figure 19 visualises the time evolution of the relative reconstruction error ϵ_u in data assimilation when initialised with (i) an interpolated field (4DVar), and (ii) the super-resolved field (4DVarSR) for assimilation windows of lengths $T_{DA}/T_e \in \{0.21, 0.43\}$. The curves in this figure are ensemble averages over four independent realisations, with shaded regions indicating a \pm one standard deviation around the ensemble mean. As for the lower Reynolds number, 4DVarSR is seen

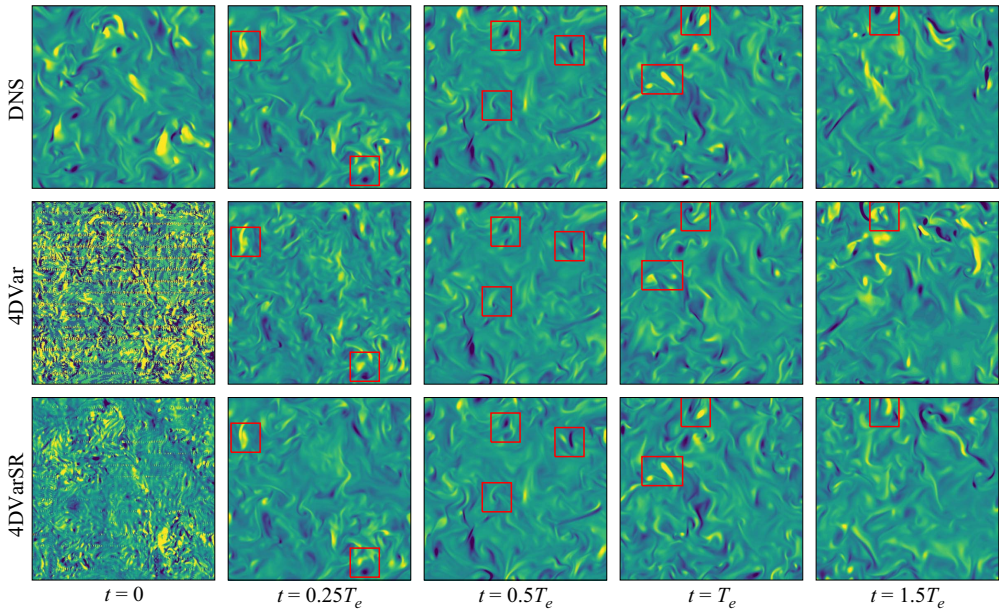


Figure 20. Comparison of the unrolled vorticity trajectory for ‘standard’ 4DVar and 4DVarSR (both $T_{DA} \approx 0.43T_e$) and a coarsening factor $M = 16$ with respect to the 256^3 grid at $Re_\lambda \approx 112$. Shown are slices of ω_y/ω_{rms} at $z = 0$, with colouring as in figure 1. Red boxes highlight individual structures in the ground truth state and their reconstructions.

to outperform data assimilation with classical interpolation over the full assimilation window, even though the network was trained on structures at the much lower $Re_\lambda \approx 70$.

In figure 20, selected snapshots of the reconstructed vorticity fields ω^p are shown for both 4DVar and 4DVarSR, as well as the target DNS trajectory. As in the lower-Reynolds-number case (see figure 14), the 4DVarSR-based reconstruction of the initial vorticity field, ω_0^p , features significantly less high-wavenumber noise than the standard 4DVar counterpart, and a comparison with the DNS ground truth state reveals that the former correctly reproduces the overall organisation of intense vorticity patterns. At $t = 0.25T_e$, the initial high-wavenumber noise in the 4DVarSR-generated fields has fully diffused and the reconstructed field is visually hard to distinguish from the corresponding ground truth state (see the red boxes in figure 20). For the 4DVar-based reconstructions, it takes most of the assimilation window length $T_{DA} \approx 0.43T_e$ until the erroneous small-scale features reminiscent of the initial high-wavenumber noise disappear (see the snapshot at $t = 0.5T_e$). Beyond the assimilation window, the 4DVar-reconstructed trajectory is seen to quickly deviate from the ground truth, with key vorticity patterns differing from those in the DNS state. On the other hand, 4DVarSR is capable of generating faithful reproductions of many features in the target snapshots even at $t = T_e$, though differences become more prominent. Over yet longer time horizons (see the rightmost snapshots at $t = 1.5T_e$), data assimilation is no longer able to provide an accurate prediction of the flow evolution, irrespective of the flow state it was initialised with. The latter is to be expected, given the level of downsampling and the much shorter assimilation windows under consideration.

5. Conclusion

In this paper we have presented a trajectory-based super-resolution technique for state estimation in triply periodic homogeneous isotropic turbulence. The training process is

inspired by variational data assimilation in that the network reconstructions are unrolled in time and trained against time series of coarse-grained velocity fields only – network training does not require a library of high-resolution snapshots (Page 2025*b*). A fully differentiable three-dimensional pseudo-spectral DNS code was used to advance the network outputs in time, which enables gradient descent on a loss involving entire trajectories.

The trained networks generate robust reconstructions of previously unseen high-resolution velocity snapshots from coarse-grained snapshots of homogeneous isotropic turbulence at $Re_\lambda \approx 70$. Notably, the initial pointwise reconstruction errors achieved by the super-resolution approach represent a 50 % or more reduction to classical 4DVar, which is an optimisation to match low-resolution observations along the specific target trajectory. Unrolling the reconstructed fields in time, 4DVar successively improves in its predictive accuracy, reaching a comparable precision to super-resolution after 0.2–0.5 large-eddy times (depending on the coarse-graining level). The good performance of super-resolution at early times can presumably be attributed to the networks having learned to construct physically realistic representations of the solution manifold from exposure to a large variety of coarse trajectories during the training. In a second step, we explored an alternative initialisation for the 4DVar algorithm using the super-resolution neural networks trained on coarse trajectories. The initialisation with a super-resolved field is highly beneficial for the performance of 4DVar, with the hybrid approach providing more accurate predictions of the ground truth trajectory over the entire assimilation window. Notably, robust predictions are achieved even for coarse-graining levels clearly above the synchronisation limit l_C at which data assimilation usually struggles.

The results showcase the versatility of the new super-resolution technique and highlight the benefits of including fundamental physical principles in the network training process. Naturally, the robust predictions achieved by trajectory-based super-resolution raise the question of how the models will perform at higher Reynolds numbers. For a moderate increase of the Reynolds number to $Re_\lambda \approx 112$, we demonstrated that the models trained at $Re_\lambda \approx 70$ are able to generate robust reconstructions of the target snapshots without re-training. The reconstructed fields again significantly improve the performance of variational assimilation when used for initialisation. When targeting significantly higher Reynolds numbers for which turbulence exhibits a pronounced inertial range over a wide range of scales, re-training will likely become necessary. However, the significant computational resources that are required for such an undertaking should not be disregarded – as a reference, training of the current networks at a target resolution of 128^3 grid points took roughly three hours per epoch on a 140 GB NVIDIA H200 card. Each training epoch involves $O(10^3)$ individual data assimilations performed in parallel over mini batches of size 16, so the full training process is equivalent to performing more than 11 000 individual assimilations. The advantage though is that once the model is trained, reconstructing a high-resolution field reduces to a single function call, while data assimilation has to be redone for each individual trajectory. Conversely, performing data assimilation starting from a super-resolved initial guess does not come with additional costs compared with a traditional initialisation once the model has been trained. Given the expected high computational costs when targeting significantly higher Reynolds numbers, it might be beneficial to switch to a model architecture that specifically takes into account the multiscale structure of the turbulent fields (Fukami, Fukagata & Taira 2019). Also, recent experience in the low-data limit might help to maintain the reconstruction accuracy when training is performed on smaller datasets (Fukami & Taira 2024).

Moreover, we see great potential in generalising the concept of a trajectory-based learning to situations for which observations are only available at arbitrarily distributed




Figure 21. Standard JPDF of (a) the second (Q) and third (R) invariant of the velocity gradient tensor and (b) the enstrophy- (Q_W) and strain-rate-related contributions (Q_S) to Q . All quantities are normalised with the spatio-temporal mean of Q_W . The blue shaded areas indicate the local value of the JPDF in the current simulations, in form of an ensemble average over 100 independent long-time trajectories. Thin black contours represent statistics from the DNS of Ooi *et al.* (1999) at $Re_\lambda \approx 71$, using a different force profile, thick black curves in (a) indicate $R = 0$ and $Q^3 + 27R^2/4 = 0$. Isocontours of the JPDFs range from 10^{-5} to 10^{-1} , with neighbouring contour lines being separated by one decade.

probing or sensor locations as in many experiments. While the ‘classical’ CNN architectures used here cannot directly deal with such irregularly distributed input data, other model architectures (e.g. graph neural networks or point-cloud-based convolutions) can easily be incorporated into the training procedure outlined in this work. There is nothing inherent in the training algorithm presented here that restricts the approach to periodic domains, or even rectilinear configurations. To generalise to more complex situations – for instance, where measurements are made on the surface of an immersed body – one would require a differentiable code capable of handling such flows, but no modification to the training loop itself. We believe our approach to be particularly attractive in these situations where one would not necessarily want to generate a large library of reference data in favour of working directly with measurements.

Acknowledgements. MW thanks the Karlsruhe House of Young Scientists (KHYS) for their financial support of his research stay at the University of Edinburgh. JP acknowledges support from a UKRI Frontier Guarantee Grant EP/Y004094/1. Computational resources were kindly provided by the Edinburgh International Data Facility (EIDF), the Data-Driven Innovation Programme at the University of Edinburgh and on the supercomputer bwUniCluster funded by the Ministry of Science, Research and the Arts Baden-Württemberg and the Universities of the State of Baden-Württemberg. The computer resources, technical expertise and assistance provided from the technical staff are gratefully acknowledged.

Declaration of interests. The authors report no conflict of interest.

Appendix A. Validation of the DNS code

To validate the spectral, three-dimensional differentiable DNS code, ensemble-averaged statistics of the long-time DNS trajectories used for the network training in § 3 are compared with DNS datasets of homogeneous isotropic turbulence from Ooi *et al.* (1999) at $Re_\lambda \approx 71$ at an identical spatial resolution $N_i = 128$, but driven by a different forcing scheme. Figure 21 shows the standard JPDFs of the second (Q) and third (R) invariant of the velocity gradient tensor (cf. figure 21a) and of the enstrophy- (Q_W) and strain-rate-related contributions (Q_S) to the second invariant (cf. figure 21b). For both measures, the

Setting	N_i	$k_{(m,N)}\eta_K$	$\Delta x_{(e)}/\eta_K$	$\Delta x_{(e)}/l_C$	Re_L	Re_λ	L/η_K	L_e/η_K	λ_T/η_K
DNS	128	1.76	1.17	0.07	862.34	75.23	150.17	158.56	17.05
$M=8$	16	0.33	9.39	0.60					
$M=16$	8	0.17	18.77	1.19					
$M=32$	4	0.08	37.54	2.39					

Table 2. Physical and numerical parameters of the DNS database for the Kolmogorov flow setting studied by Li *et al.* (2020), together with information on the low-resolution fields for a coarsening factor M . Definitions match those in table 1.

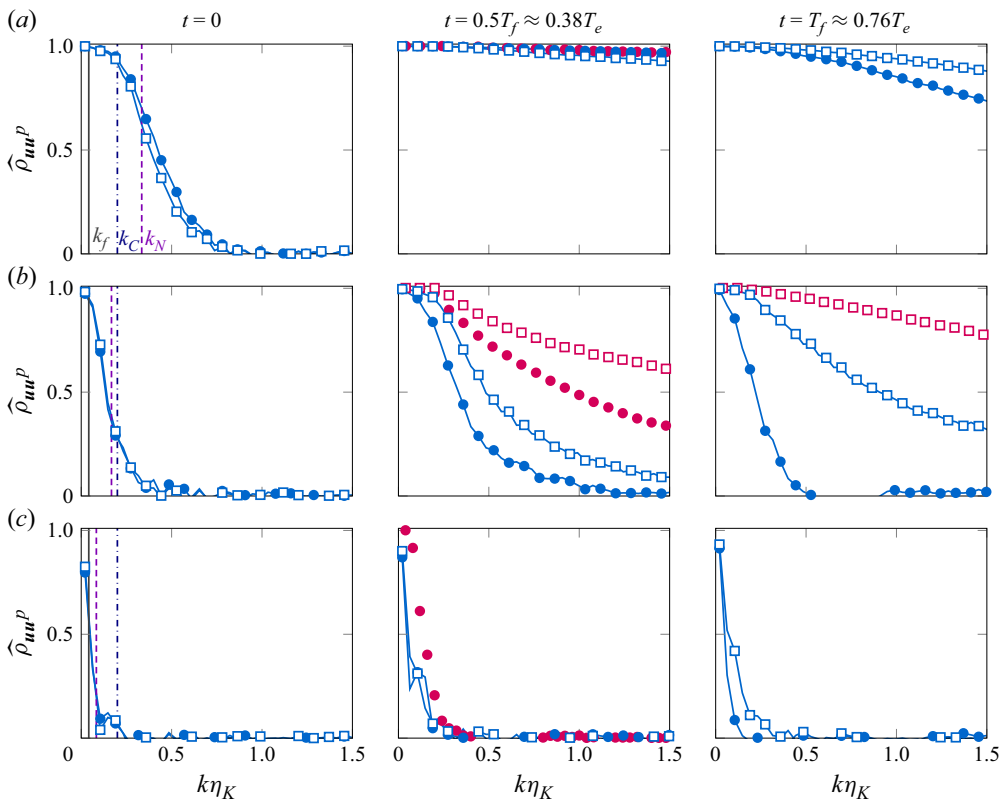


Figure 22. Normalised co-spectrum $\hat{\rho}_{uu^p}(k, t)$ between the ground truth \mathbf{u} and the 4DVar-based reconstruction \mathbf{u}^p at $t/T_f \in \{0, 0.5, 1\}$ for the Kolmogorov flow setting studied by Li *et al.* (2020); coarsening factors are (a) $M=8$ ($k_N=8$), (b) $M=16$ ($k_N=4$) and (c) $M=32$ ($k_N=2$). Results shown are the current 4DVar runs (blue) and data from Li *et al.* (2020) for $k_O=k_N$ (red) for $T_{DA}=0.5T_f$ (solid circles) and $T_{DA}=T_f$ (open squares). Note that the results of Li *et al.* (2020) are obtained as ensemble averages over $O(1)$ individual realisations and are not available for all times. Vertical lines indicate the forcing wavenumber $k_f=1$ (solid, grey), the Nyquist cutoff wavenumber k_N of the respective coarsening (dashed, purple) and the critical wavenumber $k_C=0.2\eta_K^{-1}$ (dash-dotted, blue), respectively.

current DNS runs almost perfectly collapse with the reference data, underlining that the code implementation is able to accurately capture the turbulent statistics. The remaining small deviations can presumably be explained by the much larger ensemble over which statistics have been computed in the current work.

Appendix B. 4DVar in three-dimensional Kolmogorov flow

As described in § 4.1, the 4DVar procedure proposed by Li *et al.* (2020) operates in spectral space and is therefore provided with the ground truth Fourier modes at wavenumbers $0 \leq k \leq k_O$ as the measurements for comparison. To quantify the difference in performance between this purely spectral approach and the physical-space 4DVar performed here, we present some additional physical-space 4DVar runs here that have been performed for the unidirectional Kolmogorov forcing and Reynolds number $Re_\lambda \approx 75$ to match the flow configuration considered in Li *et al.* (2020). The relevant physical and numerical parameters of these simulations are summarised in table 2. For consistency with the data presented in Li *et al.* (2020), time is henceforth measured in terms of a forcing-related time scale $T_f := u_{rms}/\chi$, where χ is the forcing amplitude in the Kolmogorov forcing

$$\mathbf{f}_e^* = (0, \chi \cos(2\pi k_f x^*/L^*), 0)^T \quad (\text{B1})$$

before non-dimensionalisation. Here, the forcing wavenumber is set at $k_f = 1$ and $T_f \approx 0.76T_e$ for the chosen parameter point.

The mode-wise reconstruction error $\hat{\rho}_{uu^p}$ is shown in figure 22 for coarse-graining levels $M \in \{8, 16, 32\}$ and two different assimilation windows with length $T_{DA}/T_f \in \{0.5, 1\}$. Where available, corresponding data points from Li *et al.* (2020) (see their figure 7c,e) with matching cutoff wavenumber for the spectral low-pass filter $k_O = k_N$ are included in red. The largest differences between the Fourier-based 4DVar results of Li *et al.* (2020) and those obtained in the current work based on coarse-grained velocity snapshots are seen at $M = 16$, where the former leads to a more accurate reproduction of the small scales. At the weakest coarse-graining $M = 8$, on the other hand, both methods lead to almost identical results. Finally, for $M = 32$, relevant deviations are restricted to the first few Fourier modes; both methods struggle to reproduce the modes at $k > 0.25/\eta_K$.

REFERENCES

- ALHASHIM, M.G., HAUSKNECHT, K. & BRENNER, M.P. 2025 Control of flow behavior in complex fluids using automatic differentiation. *Proc. Natl Acad. Sci. USA* **122** (8), e2403644122.
- BANNISTER, R.N. 2017 A review of operational methods of variational and ensemble-variational data assimilation. *Q. J. R. Meteorol. Soc.* **143** (703), 607–633.
- BARTHÉLÉMY, S., BRAJARD, J., BERTINO, L. & COUNILLON, F. 2022 Super-resolution data assimilation. *Ocean Dyn.* **72**, 661–678.
- BAYDIN, A.G., PEARLMUTTER, B.A., RADUL, A.A. & SISKIND, J.M. 2018 Automatic differentiation in machine learning: a survey. *J. Machine Learning Res.* **18** (153), 1–43.
- BEWLEY, T.R. & PROTAS, B. 2004 Skin friction and pressure: the ‘footprints’ of turbulence. *Physica D* **196** (1), 28–44.
- BOFFETTA, G. & MUSACCHIO, S. 2017 Chaos and predictability of homogeneous-isotropic turbulence. *Phys. Rev. Lett.* **119**, 054102.
- BRADBURY, J., *et al.* 2018 JAX: composable transformations of Python+NumPy programs. Available at: <http://github.com/google/jax>.
- BRENNER, M.P., ELDRIDGE, J.D. & FREUND, J.B. 2019 Perspective on machine learning for advancing fluid mechanics. *Phys. Rev. Fluids* **4**, 100501.
- BRUNTON, S.L., NOACK, B.R. & KOUMOUTSAKOS, P. 2020 Machine learning for fluid mechanics. *Annu. Rev. Fluid Mech.* **52**, 477–508.
- BUCHTA, D.A. & ZAKI, T.A. 2021 Observation-infused simulations of high-speed boundary-layer transition. *J. Fluid Mech.* **916**, A44.
- CHEVALIER, M., HEPFFNER, J., BEWLEY, T.R. & HENNINGSON, D.S. 2006 State estimation in wall-bounded flow systems. Part 2. Turbulent flows. *J. Fluid Mech.* **552**, 167–187.
- CHOLLET, F., *et al.* 2015 Keras. Available at: <https://keras.io>.
- CHONG, M.S., PERRY, A.E. & CANTWELL, B.J. 1990 A general classification of three-dimensional flow fields. *Phys. Fluids A: Fluid Dyn.* **2** (5), 765–777.

- COLBURN, C.H., CESSNA, J.B. & BEWLEY, T.R. 2011 State estimation in wall-bounded flow systems. Part 3. The ensemble Kalman filter. *J. Fluid Mech.* **682**, 289–303.
- DI LEONI, P.C., MAZZINO, A. & BIFERALE, L. 2020 Synchronization to big data: nudging the Navier–Stokes equations for data assimilation of turbulent flows. *Phys. Rev. X* **10**, 011023.
- DONG, C., LOY, C.C., HE, K. & TANG, X. 2016 Image super-resolution using deep convolutional networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **38** (2), 295–307.
- DRESDNER, G., KOCHKOV, D., NORGAARD, P., ZEPEDA-NÚÑEZ, L., SMITH, J.A., BRENNER, M.P. & HOYER, S. 2023 Learning to correct spectral methods for simulating turbulent flows. *Trans. Machine Learning Res.* 1–15.
- EVENSEN, G. 1994 Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.* **99** (C5), 10143–10162.
- EVENSEN, G. 2007 *Data Assimilation – The Ensemble Kalman Filter*, 1st edn. Springer.
- FRERIX, T., KOCHKOV, D., SMITH, J., CREMERS, D., BRENNER, M. & HOYER, S. 2021 Variational data assimilation with a learned inverse observation operator. In *Proceedings of the 38th International Conference on Machine Learning*, Proceedings of Machine Learning Research, vol. 139, pp. 3449–3458. PMLR.
- FUKAMI, K., FUKAGATA, K. & TAIRA, K. 2019 Super-resolution reconstruction of turbulent flows with machine learning. *J. Fluid Mech.* **870**, 106–120.
- FUKAMI, K., FUKAGATA, K. & TAIRA, K. 2021 Machine-learning-based spatio-temporal super resolution reconstruction of turbulent flows. *J. Fluid Mech.* **909**, A9.
- FUKAMI, K., FUKAGATA, K. & TAIRA, K. 2023 Super-resolution analysis via machine learning: a survey for fluid flows. *Theor. Comput. Fluid Dyn.* **37**, 421–444.
- FUKAMI, K., GOTO, S. & TAIRA, K. 2024 Data-driven nonlinear turbulent flow scaling with Buckingham Pi variables. *J. Fluid Mech.* **984**, R4.
- FUKAMI, K. & TAIRA, K. 2024 Single-snapshot machine learning for super-resolution of turbulence. *J. Fluid Mech.* **1001**, A32.
- GRONSKIS, A., HEITZ, D. & MÉMIN, E. 2013 Inflow and initial conditions for direct numerical simulation based on adjoint data assimilation. *J. Comput. Phys.* **242**, 480–497.
- HAYASE, T. 2015 Numerical simulation of real-world flows. *Fluid Dyn. Res.* **47** (5), 051201.
- HE, K., ZHANG, X., REN, S. & SUN, J. 2016 Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778. IEEE.
- HENDRYCKS, D. & GIMPEL, K. 2016 Gaussian error linear units (GELUS), *CoRR* abs/1606.08415, [arXiv:1606.08415](https://arxiv.org/abs/1606.08415).
- HÖPFFNER, J., CHEVALIER, M., BEWLEY, T.R. & HENNINGSON, D.S. 2005 State estimation in wall-bounded flow systems. Part 1. Perturbed laminar flows. *J. Fluid Mech.* **534**, 263–294.
- HUNT, J.C.R., WRAY, A.A. & MOIN, P. 1988 Eddies, streams, and convergence zones in turbulent flows. In *CTR Proceedings of the Summer Program 1988*, pp. 193–208. Stanford University.
- INUBUSHI, M. & CAULFIELD, C.P. 2026 Synchronisation in two-dimensional damped-driven Navier–Stokes turbulence: insights from data assimilation and Lyapunov analysis. *J. Fluid Mech.* **1027**, A41.
- KALMAN, R.E. 1960 A new approach to linear filtering and prediction problems. *J. Basic Engng* **82** (1), 35–45.
- KALNAY, E. 2002 *Atmospheric Modeling, Data Assimilation and Predictability*, 2nd edn. Cambridge University Press.
- KATO, H. & OBAYASHI, S. 2013 Approach for uncertainty of turbulence modeling based on data assimilation technique. *Comput. Fluids* **85**, 2–7.
- KATO, H., YOSHIZAWA, A., UENO, G. & OBAYASHI, S. 2015 A data assimilation methodology for reconstructing turbulent flows around aircraft. *J. Comput. Phys.* **283**, 559–581.
- KELSHAW, D., RIGAS, G. & MAGRI, L. 2022 Physics-informed CNNs for super-resolution of sparse observations on dynamical systems. In *NeurIPS 2022, Machine Learning and the Physical Sciences Workshop*. Curran Associates.
- KIM, H., KIM, J., WON, S. & LEE, C. 2021 Unsupervised deep learning for super-resolution reconstruction of turbulence. *J. Fluid Mech.* **910**, A29.
- KIM, J., MOIN, P. & MOSER, R. 1987 Turbulence statistics in fully developed channel flow at low Reynolds number. *J. Fluid Mech.* **177**, 133–166.
- KINGMA, D.P. & BA, J.L. 2015 Adam: a method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May (ed. Y. Bengio & Y. LeCun)*, May 7–9. Conference Track Proceedings.
- KOCHKOV, D., SMITH, J.A., ALIEVA, A., WANG, Q., BRENNER, M.P. & HOYER, S. 2021 Machine learning–accelerated computational fluid dynamics. *Proc. Natl Acad. Sci. USA* **118** (21), e2101784118.

- LALESCU, C.C., MENEVEAU, C. & EYINK, G.L. 2013 Synchronization of chaos in fully developed turbulence. *Phys. Rev. Lett.* **110**, 084102.
- LI, Y., ZHANG, J., DONG, G. & ABDULLAH, N.S. 2020 Small-scale reconstruction in three-dimensional Kolmogorov flows using four-dimensional variational data assimilation. *J. Fluid Mech.* **885**, A9.
- LINKMANN, M. 2018 Effects of helicity on dissipation in homogeneous box turbulence. *J. Fluid Mech.* **856**, 79–102.
- MCKAY, M.E., LINKMANN, M., CLARK, D., CHALUPA, A.A. & BERERA, A. 2017 Comparison of forcing functions in magnetohydrodynamics. *Phys. Rev. Fluids* **2**, 114604.
- MO, Y. & MAGRI, L. 2025 Reconstruction of three-dimensional turbulent flows from sparse and noisy planar measurements: a weight-sharing neural network approach. [arXiv:2509.18687](https://arxiv.org/abs/2509.18687).
- MONS, V., CHASSAING, J.-C. & SAGAUT, P. 2017 Optimal sensor placement for variational data assimilation of unsteady flows past a rotationally oscillating cylinder. *J. Fluid Mech.* **823**, 230–277.
- OOI, A., MARTÍN, J., SORIA, J. & CHONG, M.S. 1999 A study of the evolution and characteristics of the invariants of the velocity-gradient tensor in isotropic turbulence. *J. Fluid Mech.* **381**, 141–174.
- PAGE, J. 2025a Computation of simple invariant solutions in fluid turbulence with the aid of deep learning. *Nonlinear Dyn.* **113**, 31895–31915.
- PAGE, J. 2025b Super-resolution of turbulence with dynamics in the loss. *J. Fluid Mech.* **1002**, R3.
- PAGE, J., NORGAARD, P., BRENNER, M.P. & KERSWELL, R.R. 2024 Recurrent flow patterns as a basis for two-dimensional turbulence: predicting statistics from structures. *Proc. Natl Acad. Sci. USA* **121** (23), e2320007121.
- POPE, S.B. 2000 *Turbulent Flows*, 5th edn. Cambridge University Press.
- SCHUMACHER, J., SCHEEL, J.D., KRASNOV, D., DONZIS, D.A., YAKHOT, V. & SREENIVASAN, K.R. 2014 Small-scale universality in fluid turbulence. *Proc. Natl Acad. Sci. USA* **111** (30), 10961–10965.
- TAIRA, K., RIGAS, G. & FUKAMI, K. 2025 Machine learning in fluid dynamics: a critical assessment. *Phys. Rev. Fluids* **10**, 090701.
- TALAGRAND, O. 1997 Assimilation of observations, an introduction. *J. Met. Soc. Japan Ser. II* **75** (1B), 191–209.
- TALAGRAND, O. & COURTIER, P. 1987 Variational assimilation of meteorological observations with the adjoint vorticity equation. I: theory. *Q. J. R. Meteorol. Soc.* **113** (478), 1311–1328.
- UM, K., BRAND, R., FEI, Y.R., HOLL, P. & THUREY, N. 2020 Solver-in-the-loop: learning from differentiable physics to interact with iterative PDE-solvers. In *Advances in Neural Information Processing Systems* (ed. H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan & H. Lin), vol. 33, pp. 6111–6122. Curran Associates, Inc.
- VELA-MARTÍN, A. 2021 The synchronisation of intense vorticity in isotropic turbulence. *J. Fluid Mech.* **913**, R8.
- WANG, M., WANG, Q. & ZAKI, T.A. 2019 Discrete adjoint of fractional-step incompressible Navier–Stokes solver in curvilinear coordinates and application to data assimilation. *J. Comput. Phys.* **396**, 427–450.
- WANG, M. & ZAKI, T.A. 2021 State estimation in turbulent channel flow from limited observations. *J. Fluid Mech.* **917**, A9.
- WANG, M. & ZAKI, T.A. 2022 Synchronization of turbulence in channel flow. *J. Fluid Mech.* **943**, A4.
- WANG, M. & ZAKI, T.A. 2025 Variational data assimilation in wall turbulence: from outer observations to wall stress and pressure. *J. Fluid Mech.* **1008**, A26.
- WANG, Q., WANG, M. & ZAKI, T.A. 2022 What is observable from wall data in turbulent channel flow? *J. Fluid Mech.* **941**, A48.
- YASUDA, Y. & ONISHI, R. 2023 Spatio-temporal super-resolution data assimilation (SRDA) utilizing deep neural networks with domain generalization. *J. Adv. Model. Earth Sys.* **15** (11), e2023MS003658.
- YOSHIDA, K., YAMAGUCHI, J. & KANEDA, Y. 2005 Regeneration of small eddies by data assimilation in turbulence. *Phys. Rev. Lett.* **94**, 014501.
- ZAKI, T.A. 2025 Turbulence from an observer perspective. *Annu. Rev. Fluid Mech.* **57**, 311–334.