



Dynamical relevance of periodic orbits under increasing Reynolds number and connections to inviscid dynamics

Andrew Cleary¹  and Jacob Page¹ 

¹School of Mathematics and Maxwell Institute for Mathematical Sciences, University of Edinburgh, Edinburgh EH9 3FD, UK

Corresponding author: Jacob Page, jacob.page@ed.ac.uk

(Received 11 February 2025; revised 3 September 2025; accepted 4 September 2025)

Large numbers of relative periodic orbits (RPOs) have been found recently in doubly periodic, two-dimensional Kolmogorov flow at moderate Reynolds numbers $Re \in \{40, 100\}$. While these solutions lead to robust statistical reconstructions at the Re values where they were obtained, it is unclear how their dynamical importance changes with Re . Arclength continuation on this library of solutions reveals that large numbers of RPOs quickly become dynamically irrelevant, reaching dissipation values either much larger or smaller than the values typical of the turbulent attractor at high Re . The scaling of the high-dissipation RPOs is shown to be consistent with a direct connection to solutions of the unforced Euler equation, and is observed for a wide variety of states beyond the ‘unimodal’ solutions considered in previous work (Kim & Okamoto, *Nonlinearity* vol. 28, 2015, p. 3219). However, the weakly dissipative states have properties indicating a connection to exact solutions of a forced Euler equation. The dynamical irrelevance of many solutions leads to poor statistical reconstruction at higher Re , raising serious questions for the future use of RPOs for estimating probability densities. Motivated by the Euler connection of some of our RPOs, we also show that many of these states can be well described by exact relative periodic solutions in a system of point vortices. The point vortex RPOs are converged via gradient-based optimisation of a scalar loss function which (i) matches the dynamics of the point vortices to the turbulent vortex cores and (ii) insists the point vortex evolution is itself time-periodic.

Key words: nonlinear dynamical systems, vortex dynamics

1. Introduction

The dynamical systems view of turbulence considers turbulent flows as trajectories in a high dimensional space, pin-balling between simple invariant solutions, such as equilibria, travelling waves (TWs), unstable periodic orbits (UPOs) or relative periodic orbits (RPOs) in systems with continuous translational symmetry (Kawahara, Uhlmann & van Veen 2012; Graham & Floryan 2021). This perspective gained considerable interest in recent years following the discovery of unstable travelling waves on the laminar–turbulent boundary in a pipe (Faisst & Eckhardt 2003; Wedin & Kerswell 2004) and a UPO in a minimal Couette flow (Kawahara & Kida 2001). Since then, there have been systematic searches for these simple invariant solutions in a myriad of flow configurations including pipe flow (Kerswell 2005; Budanur *et al.* 2017), planar channels (Waleffe 2001; Park & Graham 2015), plane Couette flow (Gibson, Halcrow & Cvitanović 2008; Halcrow *et al.* 2009), Taylor–Couette flow (Krygier, Pughe-Sanford & Grigoriev 2021), body-forced two-dimensional flows in periodic (Chandler & Kerswell 2013; Lucas & Kerswell 2015; Zhigunov & Grigoriev 2023; Page *et al.* 2024*b*) and wall-bounded (Suri *et al.* 2020) configurations, and three-dimensional Kolmogorov flows (Yalnız *et al.* 2021; Lucas & Kerswell 2017), to list only a few.

There have been promising instantaneous observations of apparent visits to unstable simple invariant solutions in experiments and simulations (Hof *et al.* 2004; Suri *et al.* 2017, 2018), with perhaps the most convincing being the Taylor–Couette study by Krygier *et al.* (2021) showing clear evidence of shadowing of RPOs by the turbulence over multiple periods. This observed shadowing is consistent with the view in periodic orbit theory, which was developed to predict statistics of chaotic attractors in uniformly hyperbolic dynamical systems (Artuso *et al.* 1990*a,b*; Cvitanović 1991; Cvitanović *et al.* 2016). Periodic orbit theory expresses statistics as weighted sums of the statistics of UPOs, where the weights are determined by the local stability of each solution (Christiansen, Cvitanović & Putkaradze 1997). These results can be straightforwardly adapted to RPOs in systems with continuous symmetry (Budanur, Borrero-Echeverry & Cvitanović 2015). However, rigorous application of these approaches requires a complete library of solutions (see the discussion by Cvitanović 2013) and attempts to use the formulae on a modest set of RPOs in two-dimensional turbulence proved no better than simply assigning an equal weight to each solution (Chandler & Kerswell 2013; Lucas & Kerswell 2015). This has motivated recent data-driven approaches for ‘optimal’ weight estimation within an incomplete library of exact solutions (Yalnız *et al.* 2021; Page *et al.* 2024*b*; Redfern, Lazer & Lucas 2024; Pughe-Sanford *et al.* 2025).

A long-standing limitation in the application of the above-mentioned approximate statistical reconstruction methods has been the computation of sufficiently large libraries of dynamically important simple invariant solutions. Recent attempts to remove reliance on classical recurrent flow analysis to detect RPOs have had some success in significantly increasing the number of converged solutions, leading to relatively robust statistical coverage with periodic orbits in moderate-*Re* Kolmogorov flow (Page *et al.* 2024*b*; Redfern *et al.* 2024). The hope has been that, with good statistical coverage at one value of *Re*, predictions can then be made at higher Reynolds numbers by tracking the RPOs along their solution branches, though, at present, we have only limited understanding of the range of *Re* over which we can expect robust reproduction of turbulent statistics. Previous continuation efforts have shown that solutions often become increasingly dynamically irrelevant as *Re* is increased, with many moving to the laminar–turbulent boundary in bistable flows (Waleffe 2003; Wang, Gibson & Waleffe 2007). Continuation in wall-bounded flows also has to contend with the emergence of multiscale structure not

contained in the low- Re solutions, and efforts have focused on looking for solutions with this property in reduced order models (McCormack, Cavalieri & Hwang 2024), though exact solutions have been found to describe vertically localised detached/attached eddies at high Re (Deguchi 2015; Eckhardt & Zammert 2018; Yang, Willis & Hwang 2019). Furthermore, solutions often emerge in saddle node bifurcations, while the solution branches themselves may be highly non-monotonic in Re (Gibson *et al.* 2008; Chandler & Kerswell 2013). An additional layer of complexity is the need to also understand how the weights in the statistical expansions should change under increasing Reynolds number. Here, we make a data-driven assessment of both (i) the range of statistical coverage (in Re) that is obtainable with a set of solutions converged at a single parameter setting and (ii) the dependence of the statistical weightings as Re is adjusted.

A large library of converged simple invariant solutions will describe a variety of self-sustaining processes. Systematically identifying the relevant physical phenomena is a significant challenge. In three-dimensional shear flows, asymptotic theories have emerged to describe roll/streak interactions (Hall & Smith 1991; Hall & Sherwin 2010) and the ‘staircase’-like structure of the instantaneous streamwise velocity (Montemuro *et al.* 2020). In two-dimensional turbulence, there is evidence that ‘dynamically relevant’ exact solutions connect directly to solutions of the inviscid Euler equation, which can simplify the elucidation of the underlying mechanics. This connection has been conjectured for a particular class of ‘unimodal’ solutions in two-dimensional Kolmogorov flow (Kim & Okamoto 2015; Kim, Miyaji & Okamoto 2017) which share similar vortical features with the condensate that dominates undamped two-dimensional turbulence at high Re (Smith & Yakhot 1993) and which become insensitive to the background forcing required to maintain them in the presence of viscosity. A collection of exact solutions to the Euler equation (with a regularising hyperviscous term) was obtained recently by Zhigunov & Grigoriev (2023). The authors were able to show that many of the exact Euler solutions dominated by a large vortex pair were relevant to the dynamics of finite- Re turbulence. In a later study, Reynoso, Zhigunov & Grigoriev (2024) found an inviscid mechanism – via a self-similar class of solutions to the Euler equation assuming a background large-scale hyperbolic flow – to explain the mechanism of the direct cascade.

A generic feature of the various Euler solutions reported by Zhigunov & Grigoriev (2023) is that they exist as members of infinite-dimensional continuous families. This feature is observed in other analytical (smooth) solutions to the Euler equations, with perhaps the most well known being ‘Stuart vortices’ (Stuart 1967), a row of spanwise ‘cat’s eye’ vortices with uniform counter-flowing streams in the far field. The continuous family depends on a single parameter, which smoothly deforms the flow from a parallel tanh mixing layer through the cat’s-eye structures to a row of point vortices. Analogous solution families have been obtained on the surface of a stationary sphere (Crowdy 2003) and the 2-torus (Sakajo 2019). Other inviscid solutions with some relevance to the various Navier–Stokes RPOs found in the present study include multipolar solutions consisting of a central core vortex surrounded by a number of opposite-signed satellite vortices. These structures have been found analytically as superpositions of vortex patches with line vortices (Crowdy 1999) and in numerical simulations of instability growth on axisymmetric vortices of various forms (Carton, Flierl & Polvani 1989; Carnevale & Kloosterziel 1994; Morel & Carton 1994).

Connections to exact solutions of the inviscid equations have also been conjectured in decaying two-dimensional turbulence. For instance, Jiménez (2020) showed that most of the kinetic energy in the early stages of decay, during which the dominant scale of the kinetic energy is small compared with the domain, is contained in a slowly evolving quasi-equilibrium ‘crystal’ of vortex cores. The significance of point vortex dynamics in

the description of two-dimensional turbulence is well known, with compelling numerical evidence that the dynamics of finite-area vortices can be described by a ‘punctuated Hamiltonian’ model (Benzi *et al.* 1992). In this model, vortex motion is governed by the equations of point vortex dynamics, interrupted only by non-Hamiltonian merger events. There are large numbers of vortex crystals (this terminology is reserved for exact relative equilibria with fixed inter-vortex distances, see the review by Aref *et al.* 2003) documented in the literature for various numbers of vortex cores, N_v . Many solutions have been obtained via geometrical or symmetry considerations (e.g. Stieltjes 1900; Lewis & Ratiu 1996; Aref & Buren 2005). In another approach, free-energy minimisation has been employed to find extremely large numbers of crystals for large $N_v = O(50)$ (Campbell & Ziff 1979; Cleary & Page 2023). The latter study made use of gradient based optimisation on augmented loss functions to search for crystals with specific features, and is adapted here to find exact point vortex solutions that replicate the key vortex dynamics contained in a library of Navier–Stokes RPOs.

The aims of this paper are twofold. The first is to investigate how the dynamical importance of RPOs in two-dimensional Kolmogorov flow changes with increasing Re . This is achieved by performing a large arclength continuation in Re of a library of RPOs and tracking their contribution to the turbulent statistics. One intriguing feature of the solutions obtained at relatively low Re is that many appear to connect directly to solutions of the unforced Euler equation as $Re \rightarrow \infty$. As a result, a second aim of this manuscript is to attempt to represent the dynamics of the self-sustaining, large-scale coherent vortices in the RPOs with exact solutions of the point vortex system. In § 2, we formulate the equations of motion of two-dimensional Kolmogorov flow and the point vortex system in doubly periodic domains. In § 3, we present the large continuation effort of turbulent RPOs. In § 4, we present the labelling of the turbulent RPOs via solutions of the point vortex model.

2. Formulation

We begin this section by introducing two-dimensional Kolmogorov flow and the associated library of RPOs (§ 2.1) that seed the continuation effort presented in § 3. We present a method to extract vortex cores from the RPOs in § 2.2, before formulating the doubly periodic point vortex system in § 2.3, with which we will model the large-scale vortex dynamics in the Kolmogorov solutions.

2.1. Kolmogorov flow

We consider two-dimensional turbulence in a square domain with periodic boundary conditions, driven by a monochromatic body force in the streamwise direction. The out-of-plane vorticity $\omega = \partial_x v - \partial_y u$, where the velocity $\mathbf{u} = (u, v)$, evolves according to

$$\partial_t \omega + \mathbf{u} \cdot \nabla \omega = \frac{1}{Re} \nabla^2 \omega - n \cos ny. \quad (2.1)$$

In this non-dimensionalisation, we have chosen a length scale $1/k^*$, which is the inverse of the fundamental wavenumber of the domain $k^* = 2\pi/L^*$, and a time scale $1/\sqrt{k^* \chi^*}$, where χ^* is the amplitude of the forcing in the momentum equation. These length and time scales lead to the definition of the Reynolds number in this flow $Re := \sqrt{\chi^*}/k^{*3}/\nu$, where ν is the kinematic viscosity. Throughout this work, we set the forcing wavenumber $n = 4$, as has been common in previous studies (Chandler & Kerswell 2013; Page, Brenner & Kerswell 2021, 2024b).

Equation (2.1) is equivariant under continuous shifts in the streamwise (x) direction, $\mathcal{T}_s: \omega(x, y) \rightarrow \omega(x + s, y)$, under discrete shift-reflects by a half-wavelength in y , $\mathcal{S}: \omega(x, y) \rightarrow -\omega(-x, y + \pi/4)$, and under discrete rotations by π , $\mathcal{R}: \omega(x, y) \rightarrow \omega(-x, -y)$. The existence of the continuous symmetry means that we generically expect exact solutions to be relative periodic orbits that shift a finite distance in x over one period. This is the case for the vast majority of solutions in our library.

Some key integral observables reported throughout this work are the total kinetic energy,

$$E(t) := \frac{1}{2} \langle \mathbf{u}^2 \rangle_V, \quad (2.2)$$

the total dissipation rate,

$$D := \frac{1}{Re} \langle |\nabla \mathbf{u}|^2 \rangle_V = \frac{1}{Re} \langle \omega^2 \rangle_V, \quad (2.3)$$

and the total production rate,

$$I := \langle u \sin(ny) \rangle_V, \quad (2.4)$$

where the average over the volume V is defined as

$$\langle \bullet \rangle_V := \frac{1}{(2\pi)^2} \iint \bullet \, d^2 \mathbf{x}. \quad (2.5)$$

We solve (2.1) using the spectral version of the JAX-CFD solver (Kochkov *et al.* 2021; Dresdner *et al.* 2022). At each time step in the solver, the velocity field is computed by solving the Poisson equation $\nabla^2 \psi = -\omega$, where the streamfunction ψ is related to the induced velocity components via $u = \partial_y \psi$, $v = -\partial_x \psi$. The resolution is varied depending on the value of Re : $N_x \times N_y = 128 \times 128$ when $Re \leq 150$, rising to 256×256 when $150 < Re \leq 300$ and finally 512×512 when $Re > 300$.

Our analysis begins with a large set of RPOs converged at $Re = 40$ and $Re = 100$, which are documented by Page *et al.* (2024a) and Page *et al.* (2024b) (see those papers for full details of the solutions, including periods, shifts, Floquet exponents etc.). Altogether, we begin with 174 RPOs at $Re = 40$ and 151 RPOs at $Re = 100$. To investigate how the dynamical importance of these solutions varies with Re , we perform a continuation of the solutions at $Re = 40$ and $Re = 100$ to higher Reynolds number. The continuation results in a large set of solutions in the region $Re \in (30, 1100)$, with each RPO belonging to a solution branch initialised at either $Re = 40$ or $Re = 100$. Branches can go through multiple fold bifurcations throughout the continuation, so that Re is not necessarily monotonic along the curves.

2.2. Vortex identification

Motivated by a possible connection of exact solutions in viscous Kolmogorov flow to solutions of the Euler equation as $Re \rightarrow \infty$ (Kim & Okamoto 2015; Kim *et al.* 2017; Zhigunov & Grigoriev 2023), part of this work explores the ability of point vortex solutions to capture the dynamics of the large-scale vortices in the Navier–Stokes RPOs. Note that while a Kolmogorov forcing profile is used throughout this work, previous work (Gallet & Young 2013; Kim *et al.* 2017; Zhigunov & Grigoriev 2023) indicates that the large-scale structure of the flow may detach from the forcing in (2.1) at very high Re .

To model the vortex cores of the Kolmogorov RPOs with point vortices, we must first identify and extract coherent vortices from the turbulent flows. We follow the methodology of Page *et al.* (2024a) and first compute the root-mean-square (r.m.s.) vorticity fluctuations

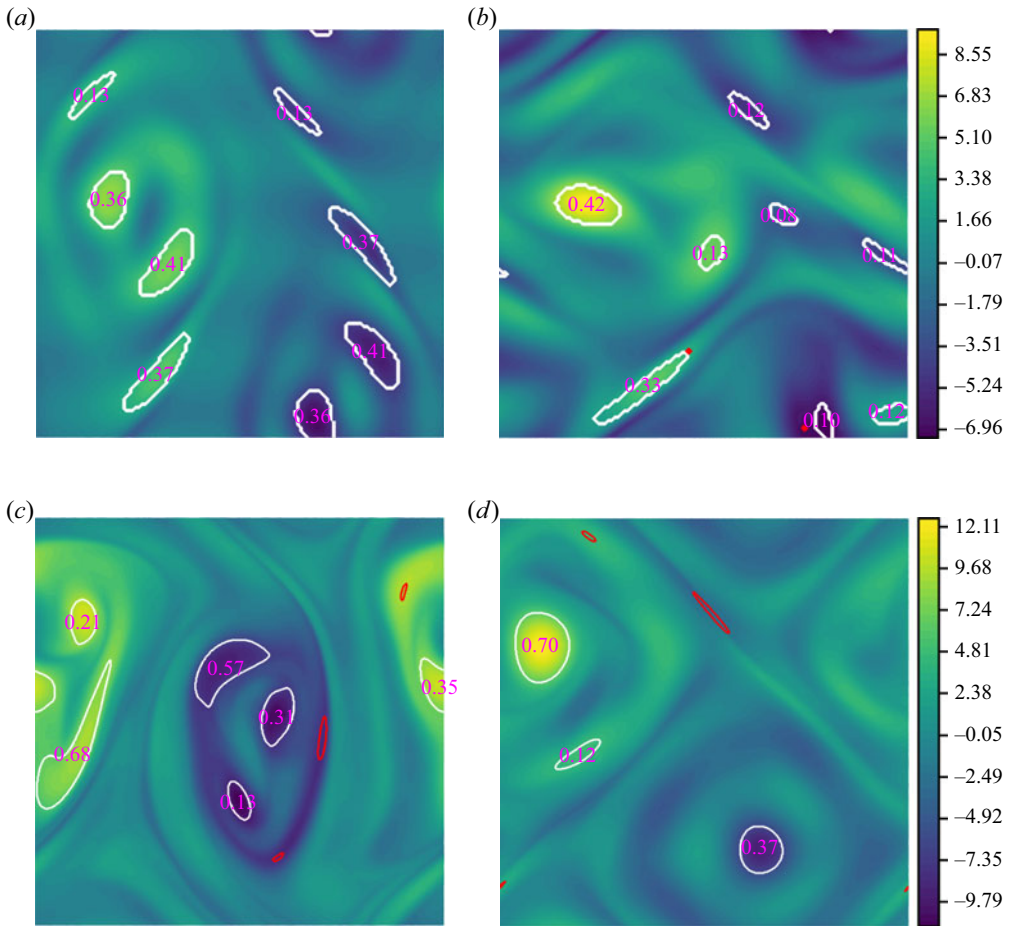


Figure 1. Application of the vortex extraction criterion discussed in the text to two example snapshots at (a,b) $Re = 100$ and (c,d) $Re = 400$. Contours are of the out-of-plane vorticity, white lines identify regions identified by the threshold (2.6), purple text indicates the region area and red lines (which here enclose areas so small they appear as points in the visualisation) highlight which of these regions were discarded due to the bound imposed on the region area.

$\omega_{rms} := \sqrt{\langle (\omega(\mathbf{x}, t) - \overline{\omega}(y))^2 \rangle_V}$, where the overline denotes an average over time and x direction. Spatially localised vortices are then extracted as any connected regions \mathcal{V}_i , where

$$|\omega(\mathbf{x}, t) - \omega_{rms}| \geq 2\omega_{rms}. \quad (2.6)$$

Some examples of these connected regions extracted from a vorticity snapshot are shown in figure 1 at both $Re = 100$ and $Re = 400$. Note that this method requires that a reference vorticity magnitude must be chosen ($2\omega_{rms}$ in (2.6)), which influences the boundaries of the extracted vortex cores. This method is agnostic to the shape of the connected region, so that strong vorticity filaments are also occasionally extracted as vortex ‘cores’ – note the thin structure stretching across the periodic boundary in figure 1(b).

Following vortex boundary identification, the circulation $\Gamma_i := \int_{\mathbf{x} \in \mathcal{V}_i} \omega(\mathbf{x}) d^2\mathbf{x}$, vortex area $A_i := \int_{\mathbf{x} \in \mathcal{V}_i} \mathbb{I}(\mathbf{x}) d^2\mathbf{x}$ and the centre of vorticity $\tilde{\mathbf{x}}_i = (\tilde{x}_i, \tilde{y}_i) = (1/\Gamma_i) \int_{\mathbf{x} \in \mathcal{V}_i} \omega(\mathbf{x})(x dx, y dy)$ of each vortex core can all be computed. Similar to the approach

by Benzi, Patarnello & Santangelo (1987), we consider only ‘large-scale’ vortices, ignoring any structures with area less than 10 % of the largest vortex $A_i < 0.1 \max_j (A_j)$ (see discarded vortices in red in figure 1). Finally, to satisfy the Gauss constraint, $\int \omega(\mathbf{x}) \, d^2\mathbf{x} = 0$, we must ensure that $\sum_j \Gamma_j = 0$, which is enforced by equally modifying the extracted circulation of each vortex core, $\Gamma_j \rightarrow \Gamma_j - \langle \Gamma \rangle$, prior to initialising a point vortex computation.

Our attempt to fit point vortex solutions is performed at $Re = 100$. At this value, we find an average $\langle |\Gamma| \rangle = 2.84$ with standard deviation $\sigma_\Gamma = 2.68$ for the extracted cores in the RPO library (the distribution has a long right tail, see Page *et al.* 2024a). As a result, the induced velocities are $O(1)$. This should be contrasted to the much weaker time average velocity, which is approximately an order of magnitude smaller (statistics are presented in § 3.5).

2.3. Doubly periodic point vortex model

We now outline the formulation for the point vortex problem in a doubly periodic domain. The following equations of motion in this system were first derived by Weiss & McWilliams (1991). An alternative derivation is presented here following the approach adopted by Aref *et al.* (2003) for a periodic strip. We begin by considering N_v point vortices in an infinite two-dimensional domain without a background velocity, where the j th vortex is located at the complex position $z_j = x_j + iy_j$. Each point vortex moves under the induced velocity from the $N_v - 1$ other vortices in the domain. The equations of motion for these point vortices are

$$\dot{\bar{z}}_j = \dot{x}_j - i\dot{y}_j = \frac{1}{2\pi i} \sum_{k=1}^{N_v}{}' \frac{\Gamma_k}{z_j - z_k}, \quad (2.7)$$

where the overbar represents the complex conjugate, the dot denotes a time derivative, Γ_k is the circulation of the k th point vortex and the prime indicates the omission of any singular terms (i.e. no self-induced velocity from $j = k$). These equations form a Hamiltonian system (e.g. see Batchelor 1967),

$$\Gamma_j \dot{\bar{z}}_j = -\frac{1}{i} \frac{\partial H_u}{\partial z_j}, \quad (2.8)$$

where the unbounded Hamiltonian is

$$H_u = -\frac{1}{4\pi} \sum_{j=1}^{N_v} \sum_{k=1}^{N_v}{}' \Gamma_j \Gamma_k \log |z_j - z_k|. \quad (2.9)$$

This result can be generalised to the bounded, doubly periodic $L \times L$ domain by introducing ghost vortices at $z_j + nL + imL$ with $m, n \in \mathbb{Z} \setminus \{0\}$, $j \in \{1, \dots, N_v\}$. The equations of motion then become

$$\dot{\bar{z}}_j = \frac{1}{2\pi i} \sum_{n,m=-\infty}^{\infty} \sum_{k=1}^{N_v}{}' \frac{\Gamma_k}{z_j - z_k - nL - imL}. \quad (2.10)$$

We can make use of the identity

$$\sum_{n=-\infty}^{\infty} \frac{1}{x - n} = \pi \cot(\pi x) \quad (2.11)$$

to simplify (2.10) by evaluating either the summation over n

$$\dot{\bar{z}}_j = \frac{1}{2Li} \sum_k \Gamma_k \sum_{m=-\infty}^{\infty} \cot \left(\pi \left(\frac{z_j - z_k}{L} - im \right) \right), \quad (2.12)$$

or over m

$$\dot{\bar{z}}_j = \frac{1}{2L} \sum_k \Gamma_k \sum_{n=-\infty}^{\infty} \cot \left(\pi i \left(\frac{z_j - z_k}{L} - n \right) \right). \quad (2.13)$$

Equations (2.12) and (2.13) correspond physically to an infinite number of strips, periodic in the horizontal and vertical directions, respectively (Aref *et al.* 2003). We can then use the trigonometric identity

$$\cot(a + ib) = \frac{\sin(2a) - i \sinh(2b)}{\cosh(2b) - \cos(2a)} \quad (2.14)$$

to extract \dot{x}_j from (2.13) and \dot{y}_j from (2.12). This results in the final equations of motion for the doubly periodic domain

$$\begin{pmatrix} \dot{x}_j \\ \dot{y}_j \end{pmatrix} = \frac{1}{2L} \sum_{k=1}^{N_v} \Gamma_k \sum_{n=-\infty}^{\infty} \begin{pmatrix} -S_n \left(\frac{2\pi}{L}(y_j - y_k), \frac{2\pi}{L}(x_j - x_k) \right) \\ S_n \left(\frac{2\pi}{L}(x_j - x_k), \frac{2\pi}{L}(y_j - y_k) \right) \end{pmatrix}, \quad (2.15)$$

where

$$S_n(a, b) = \frac{\sin(a)}{\cosh(b - 2\pi n) - \cos(a)}. \quad (2.16)$$

The reason for extracting \dot{x}_j from (2.13) rather than (2.12) is due to the need to numerically truncate the summations. Consider, for example, the equilibrium configuration of $N_v = 4$ point vortices with equal circulation magnitude, equispaced in the periodic domain. Truncating the summation in (2.13) at finite n results in two unmatched ghost vortices in the final vertically periodic strip in the summation, which causes a very small, but non-zero induced velocity in both the x and y directions. However, this induced numerical ‘drift’ in x is substantially weaker than in y . The same argument can be made for extracting \dot{y}_j from (2.12), where the number of horizontal strips is truncated.

These equations of motion with $L = 2\pi$ match those derived via the method of Laplace transforms by Weiss & McWilliams (1991) (up to a rescaling of time to match their conventions). Weiss & McWilliams (1991) also derived the Hamiltonian for the doubly periodic system of point vortices

$$H = -\frac{1}{4\pi} \sum_{j,k=1}^{N_v} \frac{\Gamma_j \Gamma_k}{2} h \left(\frac{2\pi}{L}(x_j - x_k), \frac{2\pi}{L}(y_j - y_k) \right), \quad (2.17)$$

where

$$h(a, b) = \sum_{m=-\infty}^{\infty} \log \left(\frac{\cosh(a - 2\pi m) - \cos(b)}{\cosh(2\pi m)} \right) - \frac{a^2}{2\pi}. \quad (2.18)$$

This Hamiltonian is clearly even in both x and y . It is clearly periodic in y – periodicity in x can be shown by manipulation of the sum truncated at large M , before taking the $M \rightarrow \infty$ limit. The periodic boundary conditions break the rotational symmetry of the

unbounded domain, leaving two additional constants of motion,

$$P_x = \sum_{j=1}^{N_v} \Gamma_j x_j, \quad P_y = \sum_{j=1}^{N_v} \Gamma_j y_j, \quad (2.19)$$

arising from the translational symmetry in x and y , respectively.

We solve (2.15) using a differentiable point vortex solver JAX-PV, developed in our previous work (Cleary & Page 2023). This numerical solver is built on the JAX library (Bradbury *et al.* 2018), to allow for efficient computation of the gradient of the time-forward map $\mathbf{f}^t(\mathbf{x})$ of (2.15), where $\mathbf{x} := (x_1, y_1, \dots, x_{N_v}, y_{N_v})$, with respect to \mathbf{x} . The JAX framework also brings efficiency benefits such as just-in-time compilation and auto-vectorisation. Time integration is performed with the symplectic second order Runge–Kutta scheme at a fixed time step of $\delta t = 10^{-3}$.

3. Continuation of Kolmogorov RPOs

3.1. Continuation in Re

In this section, an arclength continuation effort of the library of RPOs in two-dimensional, doubly periodic Kolmogorov flow is presented. As discussed in § 2.1, the starting library of RPOs for the continuation comprises 174 solutions at $Re = 40$ and 151 solutions at $Re = 100$. Each of these solutions is defined by the state vector $\mathbf{X} = [\omega, s, T]^T$, where ω is a vorticity snapshot along the RPO, s is the required translational shift and T is the period of the RPO. To begin the continuation of a branch, a starting solution $\mathbf{X}(Re_0)$ at Re_0 is naively perturbed in Re by δRe , $Re_1 = Re_0 + \delta Re$. Using $\mathbf{X}(Re_0)$ as the initialisation at Re_1 , this solution is converged via a standard Newton-GMRES-Hookstep solver to yield $\mathbf{X}(Re_1)$. This is robust as long as δRe is small, but cannot work when there are turning points in the solution branch. For this reason, the continuation proceeds via arclength continuation after the first successful convergence. The branch is then instead parametrised by its arclength r , which increases monotonically along the branch. The state vector is extended to include Re , i.e. $\mathbf{X}(r) = [\omega, s, T, Re]^T$, along with an additional constraint

$$\frac{\partial \mathbf{X}}{\partial r} \cdot \frac{\partial \mathbf{X}}{\partial r} = 1, \quad (3.1)$$

to match the additional unknown. The next step size in arclength δr is controlled by

$$\delta r_{i-1} = \sqrt{(\mathbf{X}(r_{i-1}) - \mathbf{X}(r_{i-2}))^2}, \quad (3.2)$$

so that $\mathbf{X}(r_i = r_{i-1} + \delta r_{i-1})$ is computed via the modified Newton-GMRES-Hookstep solver with the extra constraint (3.1). The initialisation for $\mathbf{X}(r_i)$ was set by linearly extrapolating along r from the previous two states along the branch $\mathbf{X}(r_{i-2})$ and $\mathbf{X}(r_{i-1})$. The methodology matches that presented by Chandler & Kerswell (2013), and full details can be found in the appendix of that paper.

The continuation algorithm was automated in the following ways. For each branch, the initial perturbation in Re was set to $\delta Re = 2.5$. If the perturbed solution was not converged by the standard Newton-GMRES-Hookstep solver, δRe was halved repeatedly until a convergence was found, up to a maximum of four halvings. After this first successful step in Re , the step size in arclength δr was set according to (3.2), and $\mathbf{X}(r_i)$ was initialised using the two existing states along the branch. If the initialised state was not converged by the modified Newton-GMRES-Hookstep solver, then δr was halved repeatedly until a convergence was found, again up to a maximum of four halvings. The number of halvings

of δr required for convergence was recorded along the branch, as a proxy for the ease of convergence along the branch. If both previous solutions along the branch did not require any halvings, then δr was doubled, to speed up the continuation. Convergence of a solution was deemed to have failed if either (1) more than 100 Newtons iterations were required or (2) more than 10 Hooksteps were required more than three times. Continuation of a branch was also terminated if either (1) more than 50 solutions were converged along that branch, a restriction added due to many curves repeatedly folding, or (2) δr was halved more than four times for the current solution. The starting $Re = 40$ RPOs were continued upwards in Re , while the starting $Re = 100$ RPOs were continued in separate computations both upwards and downwards in Re .

Long trajectories of Kolmogorov flow were also simulated at various $Re \in [30, 1000]$ to compute reference turbulence statistics. We use the probability density function (p.d.f.) of the dissipation in these long computations as a simple measure of the dynamical significance of the RPOs to the turbulence, labelling a solution as ‘dynamically relevant’ if its mean dissipation rate D lies between the 1st and 99th percentile of the reference p.d.f. This method is chosen for its simplicity and ease of evaluation, but will lead to some mislabelling. Better observables should be considered in future work (see e.g. the strategies of Krygier *et al.* (2021) and Page *et al.* (2024b)). Each long trajectory was simulated for 10^5 advective time units, and sampled at intervals of 1 advective time unit. Contour plots of the p.d.f.s of the dissipation of these trajectories are shown in various figures throughout the paper to aid identification of the turbulent attractor, with boundaries of the dynamically important region of the turbulent attractor indicated by dashed black lines.

The time-averaged dissipation rate for the RPOs along each branch was computed and is plotted in figure 2 over the background turbulent p.d.f.s. This quantity was chosen to reveal the scaling of the volume-averaged velocity gradients with Re for each branch. The scaling of the time-and-volume-averaged dissipation, scaled by Re , $\overline{Re D}$, of the fully turbulent flow is also shown in the bottom panel of figure 2 and matches the ‘asymptotic’ scaling of $Re^{1/2}$ identified by Chandler & Kerswell (2013) (their observations held for $Re \leq 200$). The majority of branches seeded from RPOs at $Re = 100$ do not leave the region $Re \in [60, 200]$, with the branches repeatedly turning (this was the reason for terminating many continuations at 50 steps).

In the top panel of figure 2, the number of monotonic-in- Re subsections of solution branches is tracked – this is a proxy for the number of unique RPOs in a particular Re interval. The number of dynamically important solutions decreases more sharply with Re than the total number of monotonic subsections. This departure of the RPOs from ‘dynamically relevant’ dissipation values could be attributed to a number of causes. One is the (unlikely) exit from the attractor of RPOs via boundary crises. Alternatively, the RPOs may remain in the attractor but with an increasingly rare probability of visits from a turbulent orbit, hence, it is deemed dynamically irrelevant under our scalar-observable based criterion. Finally, there is the possibility that the solutions were not in the attractor when they converged, but close by. This is not something that is obvious from low-dimensional projections which are typically used to assess the closeness of an RPO to the turbulence, but does seem to be borne out by the following analysis.

To facilitate a discussion about the $Re \rightarrow \infty$ behaviour of the solutions, the branches are labelled as belonging to one of three classes according to the dynamical importance of the final RPO on each branch, as assessed by comparison to the turbulent PDF described previously. The final RPO of branches in class 1 have a larger dissipation rate than the dynamically important region, class 2 solutions lie in the dynamically important region (many remain confined to a small range of Re) and class 3 solutions have a smaller




Figure 2. Continuation of periodic orbits and their overlap with the turbulent dissipation p.d.f. (Top) Histograms of the number of monotonic-in- Re subsections N_{upo} of all solution branches (grey) and monotonic-in- Re subsections fully within the dynamically important region (red) as a function of Re . Both histograms are computed using 50 bins, spaced logarithmically over the range $Re \in [20, 1000]$. (Middle) The time-averaged dissipation rate $Re D$ against Re of the arclength continuation of the initial library of RPOs starting at both $Re = 40$ and $Re = 100$. The contour plot shows the reference p.d.f.s of dissipation rate, with the 1st and 99th percentiles indicated by the dashed black lines. The red/blue/green branches indicate those with terminal RPOs above/within/below this dynamically important region. Circles denote the terminal solution along each branch. Filled circles denote the branches which were terminated as 50 states were converged, while empty circles denote the branches which could not be continued further. (Bottom) The time-averaged, scaled dissipation rate $Re \overline{D}$ of a long-time simulation is shown as a function of Re (black), as well as the scaling law $Re^{1/2}$ (blue).

dissipation rate than the dynamically important region. These three classes are coloured in red, blue and green, respectively, in figure 2. At high Re , the average dissipation rate of RPOs in class 2 scales similarly to the turbulent attractor, while those in class 1 and class 3 scale more strongly (most with a scaling $D \sim Re$) or weakly ($D \sim 1/Re$ – though we are lacking large numbers of solutions in this regime), respectively. Note that the vast majority

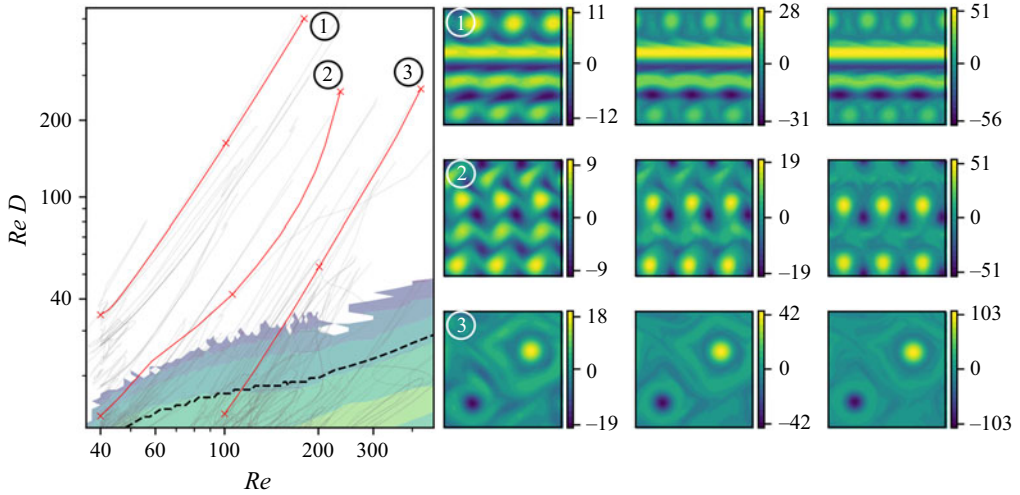


Figure 3. Time-averaged, scaled dissipation rate $Re D$ against Re for three representative branches in class 1. A single vorticity snapshot for each highlighted RPO (circled numbers 1–3) at points indicated by crosses on the branch are shown in a row on the right of the figure, with arclength increasing from left to right. Note the different colour bars for each panel, highlighting the strengthening vortical structures.

of solutions in classes 1 and 3 appear to continue monotonically with Re after they leave the dynamically relevant region.

3.2. Class 1: connections to exact unforced Euler solutions

Some example branches from solution class 1 are shown in more detail in [figure 3](#). The RPOs in this class have stronger dissipation rates than the turbulence, or equivalently larger enstrophy $\langle \omega^2 \rangle_V$. As a result, class 1 solutions have very defined vortical structures – e.g. see the horizontal bars and clear vortex cores in [figure 3](#). The strength of the local vorticity associated with these features increases dramatically as Re is raised. For example, the maximum vorticity nearly triples for the last two points highlighted in the second row of [figure 3](#), while Re increases from approximately 100 to 220.

The increasingly strong vorticity fields in class 1 meant that the baseline resolution was often inadequate as the continuation proceeded upwards in Re . As a result, some of the initial computations terminated due to a failure to converge before reaching 50 states. These branches are identified with filled diamonds in [figure 3](#), and were continued further at this point by doubling the number of Fourier modes in both spatial directions. For solutions which began at $Re = 40$, it was sometimes necessary to double the resolution again (to a maximum of 512×512) to follow the branch further upwards in Re .

The dissipation scaling $D \sim Re$ is consistent with an enstrophy $\langle \omega^2 \rangle_V \sim Re^2$, which implies $\omega = O(Re)$ if the vortex cores do not shrink. This scaling is confirmed for a number of branches in class 1 in [figure 4](#), and the velocity field (largely induced by the high amplitude vortex cores) also follows this scaling. In addition, the Re -dependence of the period of the RPOs is also shown in [figure 4](#), and scales like $T_{RPO} \propto 1/Re$. The horizontal shift associated with the class 1 RPOs ([figure 4d](#)) does not exhibit a consistent scaling with Re – though the majority solutions appear to exhibit $s_{RPO} \sim \text{constant}$ as Re increases. These scalings motivate the introduction of new variables $\Omega := \omega/Re$, $U := u/Re$ and




Figure 4. Scalings with Re of (a) the periods T_{RPO} , (b) maximum vorticity values $\max |\omega|$, (c) maximum velocity values $\max |\mathbf{u}|$ and (d) shift for a subset of branches in class 1.

$\tau := Re t$, from which (2.1) becomes

$$\partial_\tau \Omega + \mathbf{U} \cdot \nabla \Omega = \frac{1}{Re^2} (\Delta \Omega - n \cos ny). \quad (3.3)$$

An asymptotic solution would then be sought in the form of a regular perturbation series:

$$\Omega(\mathbf{x}, \tau) = \Omega_0(\mathbf{x}, \tau) + \frac{1}{Re^2} \Omega_1(\mathbf{x}, \tau) + \dots, \quad (3.4)$$

$$\mathbf{U}(\mathbf{x}, \tau) = \mathbf{U}_0(\mathbf{x}, \tau) + \frac{1}{Re^2} \mathbf{U}_1(\mathbf{x}, \tau) + \dots. \quad (3.5)$$

At each order, $\Delta \Psi_i = -\Omega_i$, where Ψ_i is the streamfunction associated with \mathbf{U}_i . The dynamics at leading order are therefore governed by the inviscid vorticity equation (the curl of the Euler equation)

$$\partial_\tau \Omega_0 + \mathbf{U}_0 \cdot \nabla \Omega_0 = 0, \quad (3.6)$$

while the first-order correction is

$$\partial_\tau \Omega_1 + \mathbf{U}_0 \cdot \nabla \Omega_1 + \mathbf{U}_1 \cdot \nabla \Omega_0 = \Delta \Omega_0 - n \cos ny. \quad (3.7)$$

The existence of a solution to the first-order problem requires that a series of solvability conditions are satisfied,

$$\int_0^T \langle \zeta_j (\Delta \Omega_0 - n \cos ny) \rangle_V d\tau = 0, \quad (3.8)$$

where – as shown by Zhigunov & Grigoriev (2023) – the $\{\zeta_j\}$ are eigenfunctions associated with continuous symmetries of the T –(relative) periodic Euler solution $\Omega_0(\mathbf{x}, \tau)$. This is a generalisation of the solvability condition of Okamoto (1994) for steady solutions, which is identical to the physical constraint derived by Batchelor (1956).

The analysis here suggests a direct connection of the finite- Re solutions to solutions of the unforced Euler equation, where it is well known that solutions exist as continuous families due to the conservation of Casimir functions (Stuart 1967; Zhigunov &

Grigoriev 2023). However, we expect that each viscous solution should connect to a single ‘member’ of an Euler family, which is selected through the viscous solvability conditions (3.8) (Okamoto 1994) – no further bifurcations associated with other ‘members’ of a solution family at finite Re are necessary as there is no direct connection.

The scalings of class 1 solutions with Re reported in figure 4 confirm the expected asymptotic behaviour. These results indicate that many Kolmogorov flow solutions found as low as $Re = 40$ – which have a structure at low Re that depends intimately on the forcing – connect directly to solutions of the unforced Euler equation. The various ‘turbulent’ RPOs found here complement the ‘unimodal’ states of Kim & Okamoto (2015), which also appear to connect to solutions of the unforced Euler equation as $Re \rightarrow \infty$. Those authors used a framework in which the body force was assumed to be $O(1/Re)$ in the momentum equation, though this is equivalent to the $O(1)$ force considered here after the rescaling used to produce (3.3); see also the discussion of Kim & Okamoto (2010). Kim & Okamoto (2015) observed unimodal Euler solutions to emerge for $500 \lesssim Re_{KO} \lesssim 100\,000$, where their Reynolds number, Re_{KO} , is related to ours by $Re \approx \sqrt{Re_{KO}}$ (approximate because of differences in the forcing profile). Their vorticity profiles, which have $\omega_{KO} = O(1)$ as $Re_{KO} \rightarrow \infty$, would also scale like $\omega = O(Re)$ using our conventions. The Re values where unimodal states are found by Kim & Okamoto (2015) correspond in our problem to a range $22 \lesssim Re \lesssim 300$, which agrees well with the Re values for which the class 1 scaling becomes apparent.

Furthermore, some of our class 1 solutions share qualitative similarities with the Euler solutions reported by Zhigunov & Grigoriev (2023). Of the 184 solutions labelled as ‘class 1’, 18 consist of a pair of opposite-signed vortices – similar profiles were reported by those authors (e.g. see example (3) in figure 3). Ratios of the time-averaged enstrophy and energy for these class 1 RPOs are typically found to be between approximately 5 and 10 (not shown), somewhat higher than those reported by Zhigunov & Grigoriev (2023), which were between 2 and 4, while the appropriately scaled periods for our RPOs are also much larger, $T Re = O(100)$. However, it is still possible that some of our states are connected to those shown by Zhigunov & Grigoriev (2023), which could presumably be checked by converging the leading (Euler) component of the solution and performing homotopy through the continuous solution families.

The class 1 solutions discussed here, some of which were highlighted in figure 3, indicate that a direct viscous-to-inviscid connection is possible for a wide variety of flow states beyond the unimodal states discussed here (Kim & Okamoto 2015; Kim *et al.* 2017; Zhigunov & Grigoriev 2023). The scaling of this solution class relative to the turbulence would seem to suggest that these states, while ‘dynamically relevant’ by our definition at the Re values where they were converged and potentially similar to turbulent trajectories, were outside of the turbulent attractor.

3.3. Classes 2 and 3

The RPOs in classes 2 and 3 tend to be dominated by two oppositely signed vortices, arranged similarly to the vortex condensates typically seen in high- Re and Euler solutions (Zhigunov & Grigoriev 2023; Page *et al.* 2024a). The large-scale opposite-signed vortex pair in class 3 branches is less distinct from the background vorticity than in class 2 branches – class 3 solutions still appear to show evidence of the $n = 4$ forcing wave. Multipolar structures (Carton *et al.* 1989; Morel & Carton 1994) are observed in the class 2 solutions. For example, note the co-rotating same-signed vortex ‘bound state’ in the top row of snapshots in figure 5, and the ‘tripole’ structures (central core with two opposite-signed satellites) in the second and third rows of figure 5. In the middle

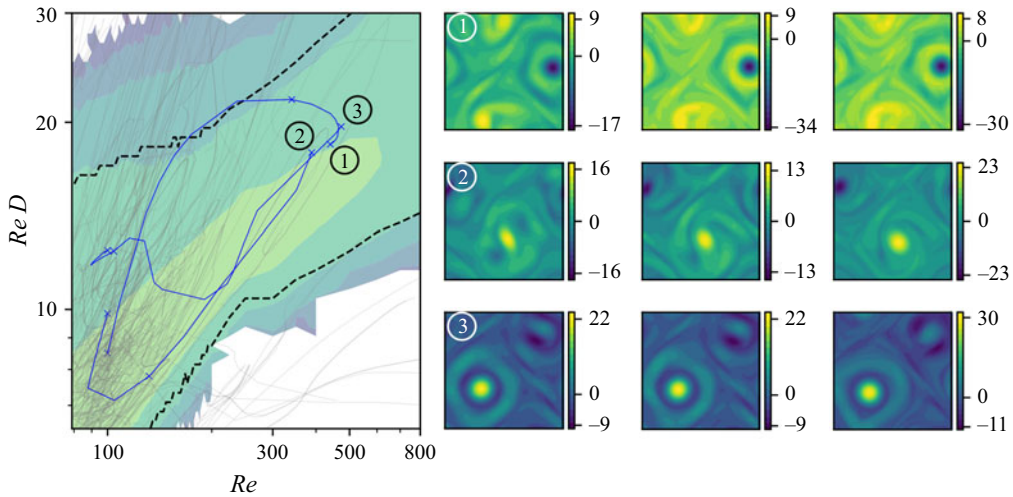


Figure 5. Time-averaged, scaled dissipation rate $Re D$ against Re for three representative branches in class 2. Three RPOs along each branch are sampled (indicated by crosses) and visualised on the right of the figure. A single vorticity snapshot for each highlighted RPO (circled numbers 1–3) at points indicated by crosses on the branch are shown in a row on the right of the figure, with arclength increasing from left to right.

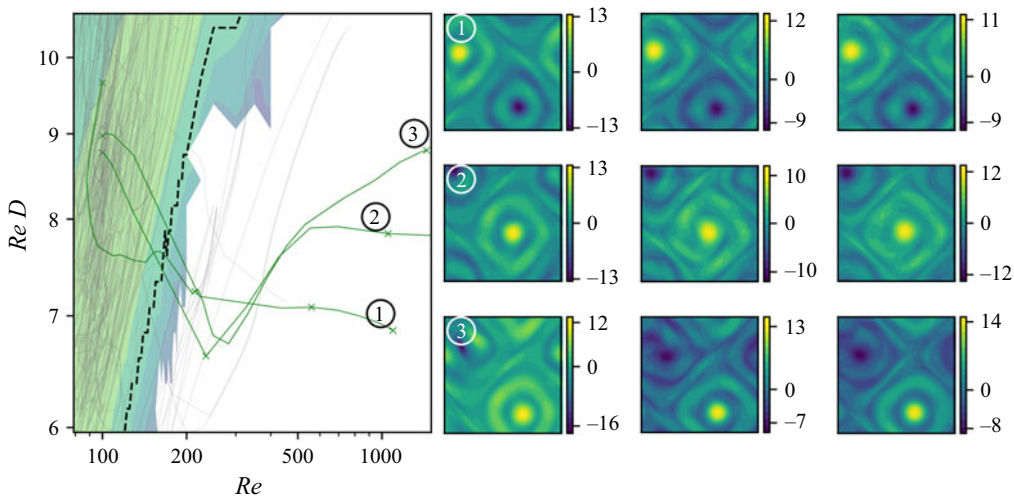


Figure 6. Time-averaged, scaled dissipation rate $Re D$ against Re for three representative branches in class 3. Three RPOs along each branch are sampled (indicated by crosses) and visualised on the right of the figure. A single vorticity snapshot for each highlighted RPO (circled numbers 1–3) at points indicated by crosses on the branch are shown in a row on the right of the figure, with arclength increasing from left to right.

row (solution ‘2’) in figure 5, the tripole structure weakens with increasing Re , and the vorticity field becomes dominated by the opposite-signed large-scale vortices. In contrast to the class 2 solutions, the vorticity strength of solutions in class 3 remains practically constant, even as Re is increased by an order of magnitude from $Re = 100$ to $Re \approx 1000$. The dominant structures visible in the solutions in class 3 also vary remarkably little with these large changes in Re .

The apparent asymptotic behaviour of the class 3 solutions, with $D \sim 1/Re$, is consistent with vorticity fields $\omega = O(1)$ (as shown in figure 6). We also observe




Figure 7. Full range of production I at each Re for the three furthest continued class 3 solutions. A logarithmic colourmap is used, and the numbering corresponds to the numbered branches in figure 6.

that $T_{RPO} \approx \text{constant}$ as Re increases (not shown). This behaviour is consistent with a connection to solutions of a *forced* Euler equation. To see this, consider a regular perturbation expansion $(\omega, \mathbf{u}) = (\omega_0, \mathbf{u}_0) + Re^{-1}(\omega_1, \mathbf{u}_1) + \dots$, from which the leading order contribution satisfies

$$\partial_t \omega_0 + \mathbf{u}_0 \cdot \nabla \omega_0 = -n \cos ny. \quad (3.9)$$

There are no dissipation mechanisms, hence, for T -periodic solutions, the forcing term must extract the same amount of kinetic energy it injects over a cycle,

$$\int_0^T I \, dt = \int_0^T \langle u_0 \sin ny \rangle_V \, dt = 0. \quad (3.10)$$

The production over a complete cycle of three class 3 RPOs is reported in figure 7 along the solution branches. For two of the branches (labelled ‘1’ and ‘3’), negative values of the production are realised at higher Re , which is consistent with a forced-Euler connection in the limit $Re \rightarrow \infty$. Further continuation would be necessary to verify this trend.

3.4. New RPOs at $Re = 100$

Due to the non-monotonicity of the solution branches, the original library at $Re = 100$ can be supplemented with new RPOs by slicing branches anywhere they cross $Re = 100$ in the arclength continuation. As a result, 101 new RPOs at $Re = 100$ were converged. The initialisations in the Newton convergence for these new RPOs were generated by linearly interpolating between the two RPOs which bracket the point where the branch crosses $Re = 100$. Some branches crossed $Re = 100$ more than once to yield more than one new unique RPO.

These new RPOs are reported in figure 8 in terms of their production, I , and dissipation D , normalised by the laminar value. Due to the localised nature in phase space of these RPOs, full coverage of the turbulent p.d.f. at this Re is difficult, and requires a large number of solutions. These new RPOs go some way to increasing the coverage of the turbulent p.d.f., though gaps remain at higher D and the tails of the I marginal p.d.f. are not covered.

3.5. Reconstruction of turbulent statistics

A variety of statistics, including full dissipation and energy p.d.f.s, were recently successfully reconstructed via an expansion in the statistics of individual RPOs at $Re = 40$ by Page *et al.* (2024b). In that paper, the weights in the expansion were determined from the invariant measure of a Markov chain, where the ‘states’ on the chain were the RPOs,




Figure 8. Energy dissipation rate D against energy production rate I at $Re = 100$, both normalised by the laminar dissipation $D_{lam} = Re/(2n^2)$. The 101 new RPOs at $Re = 100$ from the continuation are shown in red, the starting library of 151 RPOs are shown in blue. The grey background is the p.d.f. computed from a trajectory of 10^5 samples, separated by 1 advective time unit. The contour levels of the p.d.f. are spaced logarithmically. (a) New very high dissipation RPOs, far from the turbulent attractor, which resulted from the continuation of the very high dissipation RPOs at $Re = 40$. (b) Zoom-in on the turbulent attractor.

and the current state on a turbulent trajectory was assigned by measuring the distance to each RPO in the latent space of an autoencoder and identifying the closest solution. We adopt a similar approach here across $Re \in [40, 300]$, albeit without building Markov chain models. Instead, RPO weights are determined by fitting to a *single* representative statistic, and their robustness is then assessed by reconstructing other statistics using the same weights (see the approach of Redfern *et al.* 2024). We generate the various libraries of RPOs at each target Re by slicing the solution branches from the arclength continuation, as described in § 3.4. This allows us to track how the weights of each RPO change as Re is increased and assess the utility of RPOs converged at one Re value in making statistical predictions at higher values.

The number of unique solutions converged at each Re considered is shown in figure 9. There is a reduction in the number of solutions reported here compared with the number of monotonic-in- Re subsections of the solution curves reported in figure 2 due to multiple interpolated initial guesses in Newton sometimes converging to only one of the possible solutions on a branch repeatedly folding back and forth (e.g. both the lower and upper branch pair at a saddle node converging to the lower-branch solution). This could be improved by reducing the arclength step sizes in the original branch continuation. There are two distinct peaks in figure 9 at $Re = 40$ and $Re = 100$, as the arclength continuation was seeded from solutions at these values. The substantial drop in numbers of solutions at the higher Re regimes commented on in § 3.1 is also apparent – there is a near-exponential drop off in the solutions. Many of the solutions counted in figure 9 are dynamically unimportant and will not contribute to the reconstruction of the turbulent statistics.

We compute a fixed set of weights at a given Re , $\{w_j(Re)\}_{j=1}^{N_p(Re)}$, where $N_p(Re)$ is the total number of RPOs found at that Re value. The ansatz, inspired by periodic orbit theory (Artuso *et al.* 1990a; Cvitanović 2013), is that the normalised distribution p_w of any observable $\gamma : \mathcal{M} \rightarrow \mathbb{R}^n$ (here, \mathcal{M} is the state space) can be constructed via a linear




Figure 9. Number of unique RPOs converged by slicing the arclength continuation at discrete values of Re .

superposition of the RPO statistics for that same observable:

$$p_w(\gamma) = \frac{\sum_{j=1}^{N_p} w_j p_j(\gamma)}{\sum_{j=1}^{N_p} w_j}, \quad (3.11)$$

where $p_j(\gamma)$ is the distribution of the j th RPO and we are suppressing all dependencies on Re for clarity. We find the weights by minimising the Kullback–Leibler (KL) divergence between the reconstructed observable distribution and the distribution of the target turbulent statistic q ,

$$D_{KL}(p_w \| q) = \int p_w(\gamma) \log \left(\frac{p_w(\gamma)}{q(\gamma)} \right) d\gamma. \quad (3.12)$$

Optimisation of the weights proceeds via gradient-based optimisation of D_{KL} , using the Adam optimiser (Kingma & Ba 2015) with an initial learning rate of 10^{-2} . To ensure that the weights $w_j \in [0, \infty)$, we write them in terms of a ‘softplus’ function,

$$w_j = \sigma_+(\hat{w}_j) := \log(1 + e^{\hat{w}_j}) \in \mathbb{R}_+, \quad (3.13)$$

and optimise for the latent representation $\hat{w}_j \in \mathbb{R}$. The softplus function is a smooth approximation to the typical rectifier or ReLU function and ensures that the weights are strictly positive. The stopping criterion for convergence is set to

$$\frac{D_{KL}^i - D_{KL}^{i-1000}}{D_{KL}^i} < 10^{-8}, \quad (3.14)$$

where D_{KL}^i denotes the KL divergence at the i th optimiser iterate. This optimisation amounts to fitting to a one-dimensional curve and is almost instantaneous on modern hardware – negligible compared with the cost of converging the underlying solutions.

A key feature of the weights obtained via periodic orbit theory is that they are fixed by properties of the underlying exact solutions and do not vary between statistics.

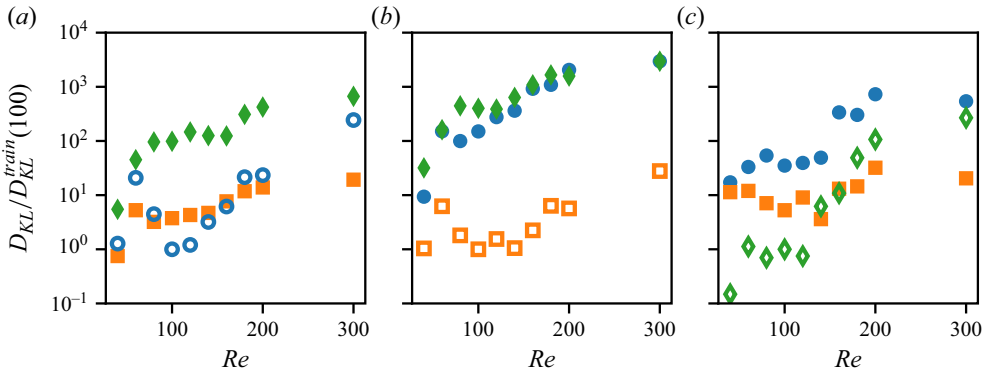


Figure 10. D_{KL} loss for the reconstruction of D/D_{lam} (blue circles), I/D_{lam} (orange squares) and E/E_{lam} (green diamonds) as a function of Re , each normalised by D_{KL} at $Re = 100$ of the training observable, which was set to D , I and E respectively in panels (a)–(c). Unfilled markers denote this training observable in each panel. The weights were then fixed when reconstructing the distributions of the other two test observables. The normalisation constants are the laminar dissipation rate $D_{lam} = Re/(2n^2)$ and the laminar energy $E_{lam} = Re^2/(4n^4)$.

Accordingly, robustness of the weights determined in the data-driven approach can be checked by computing the weights using a single ‘training’ statistic, using these fixed weights to reconstruct other ‘test’ statistics at the same Re , and then comparing KL divergences. Three observables are considered here as training statistics – the normalised dissipation rate D/D_{lam} , normalised production rate I/D_{lam} and the normalised energy E/E_{lam} .

The KL divergence of the three observables are shown as a function of Re in figure 10, setting the training observable (unfilled markers) to one of the three statistics in each panel. The KL divergences are normalised by the KL divergence of the training observable at $Re = 100$. Both D and E (figure 10a,c) as training statistics lead to relatively robust reconstruction of statistics in the range $40 \lesssim Re \lesssim 150$, with a noticeable increase in D_{KL} for all p.d.f.s beyond this point. This observation is consistent with the decreasing number of dynamically relevant solutions estimated in figure 2. The least robust observable appears to be I/D_{lam} , exhibiting the largest discrepancy between the training and test reconstructions of typically two orders of magnitude across all values of Re . The reconstruction above $Re = 300$ is not shown due to the lack of dynamically important solutions.

Unlike more standard error metrics, the KL divergence (which is not a metric) is not straightforward to interpret, and the performance of the reconstruction is assessed further in figure 11, where the p.d.f.s are compared with direct numerical simulation (DNS) ground truth, with D as the training statistic. As expected, the reconstruction of D itself is fairly well represented by the RPOs – particularly at $Re \in \{40, 100\}$ where the solutions were originally converged – but with some gaps in the higher dissipation tails. At other Re values, the reproduction of the high-dissipation tails deteriorates. This can be attributed to the emerging dynamical irrelevance of solution class 1, which climb to increasingly high dissipation rates to connect to inviscid solutions and which may be outside of the turbulent attractor. Unsurprisingly, the reconstruction of the more probable production rates is also fairly robust, though the high- I tail is only really apparent from the RPOs at $Re = 40$. However, reconstruction of the kinetic energy p.d.f. is noticeably poor compared with D and I , particularly outside of the Re where solutions were originally converged.

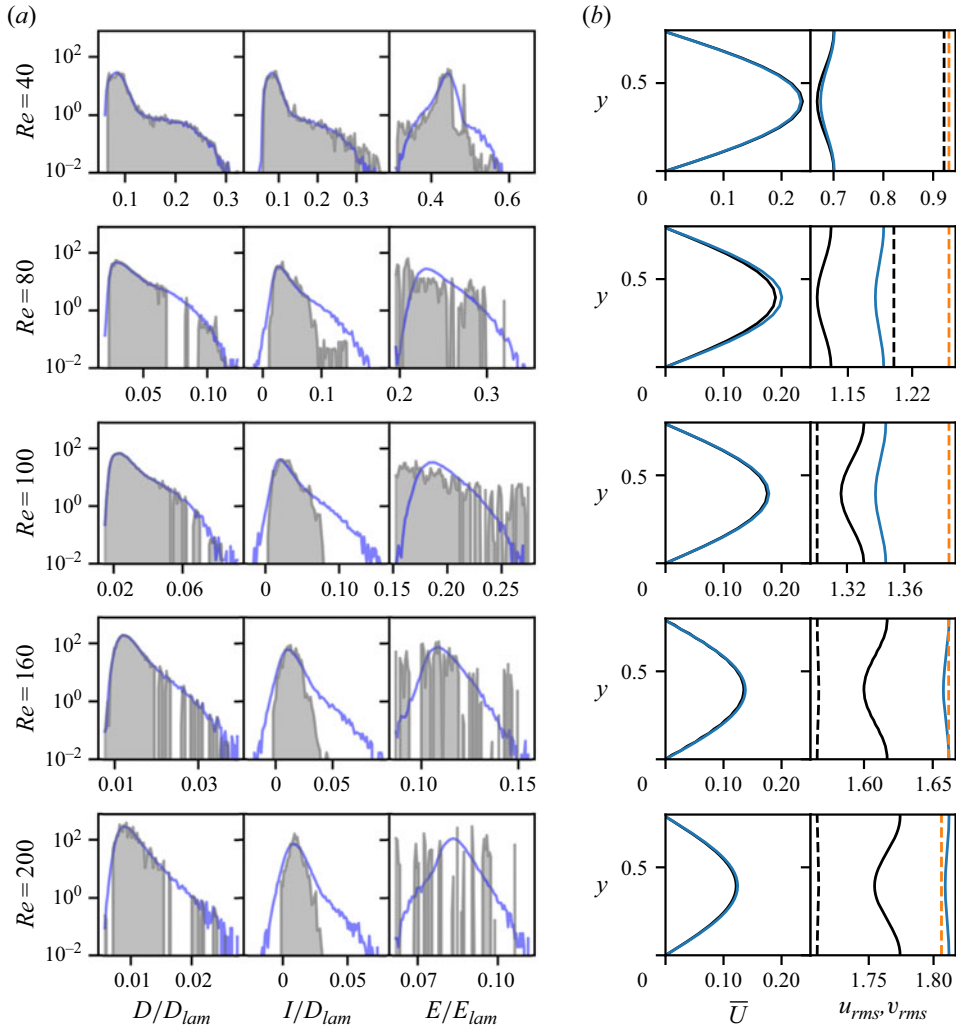


Figure 11. Reconstruction of turbulence statistics for increasing Re from the top row to the bottom row, with weights obtained using D/D_{lam} as the training observable. (a) From left to right, the reconstructed normalised dissipation rate D/D_{lam} , normalised production rate I/D_{lam} , and normalised energy E/E_{lam} are shown in grey, while the true turbulent distributions are shown in blue. (b) In the left panels, the true turbulent mean velocity profile $\bar{U}(y)$ is shown in blue, while the reconstructed profile is shown in black. In the right panels, the true turbulent r.m.s. velocity fluctuations u_{rms} (blue continuous line), v_{rms} (orange dotted line), averaged over the streamwise direction, discrete symmetries and time, are shown, while the reconstructed profiles are shown in black.

Notably, the p.d.f. reconstruction fails to represent the high- I tails or to reproduce the negative values of production observed at higher Re values. This is perhaps surprising given the relatively robust reconstruction of the dissipation and the fact that $\bar{D} = \bar{I}$ when averaging is over any relative periodic solution. However, note that the production p.d.f. is significantly wider than that of dissipation. For instance, peak values of I for $Re \gtrsim 100$ increase to more than twice the peak dissipation values (see p.d.f.s in figure 11). Therefore, a periodic orbit with time-averaged dissipation \bar{D} that samples these extreme I values will need to spend longer in the low- I region. The absence of this behaviour in our solution library is therefore likely due to the relatively short periods of the converged cycles.

Re	$\varepsilon(\hat{\mu}_D)$	$\varepsilon(\hat{\sigma}_D^2)$	$\varepsilon(\hat{\mu}_I)$	$\varepsilon(\hat{\sigma}_I^2)$	$\varepsilon(\hat{\mu}_E)$	$\varepsilon(\hat{\sigma}_E^2)$
40	0.012	0.031	0.011	0.097	0.006	0.092
60	0.089	0.022	0.09	0.056	0.025	1.576
80	0.053	0.365	0.054	0.568	0.059	1.227
100	0.013	0.136	0.015	0.563	0.037	2.812
120	0.002	0.057	0.003	0.572	0.035	2.385
140	0.01	0.138	0.012	0.602	0.008	2.469
160	0.018	0.251	0.019	0.635	0.029	2.193
180	0.025	0.358	0.03	0.697	0.0	3.857
200	0.013	0.228	0.02	0.699	0.016	6.603
300	0.029	0.295	0.171	0.32	0.003	1.656

Table 1. Relative error in predictions of the first two moments (mean $\hat{\mu}$ and variance $\hat{\sigma}^2$) in the distributions reported in [figure 11](#), where the dissipation was used as the ‘training’ observable. Relative error of a statistical estimate is defined as $\varepsilon(\hat{\mu}) := |\hat{\mu} - \mu|/\mu$, where the ground truth value (here μ) is obtained from the ‘true’ turbulent distributions shown in blue in [figure 11](#).

Prediction of a full p.d.f. is a particularly challenging task – though one common feature for all reconstructed distributions is that continuation of the RPOs from other Re values provides coverage for the more probable events. To understand the utility of the solution library at estimating point statistics, we also report in [table 1](#) the relative error in reconstruction of the first two moments of the D , I and E p.d.f.s using the RPOs. These estimates were computed from the distributions reported in [figure 11](#), and show that estimates of the mean of all three quantities is relatively robust at all Re except $Re = 300$: generally, errors of less than 5 % are observed. The variance is not so impressive, particularly for the energy, which is reflective of the fact that the solutions representing the rarer excursions to, for example, high-dissipation or high-energy regions of the state space appear to be relevant only in a small- Re range. This is true even in the ‘asymptotic’ regime beyond $Re \gtrsim 100$ – for instance, see the relative lack of solutions in the tails at $Re = 160$ in [figure 11](#) compared with those at $Re = 100$.

Similarly, predictions for the mean velocity and r.m.s. velocity fluctuations shown in [figure 11](#) (averaged over the streamwise direction, the 16 discrete symmetries of $n = 4$ Kolmogorov flow and time, see [Chandler & Kerswell 2013](#); [Farazmand 2016](#)) are also relatively robust over the range of Re considered. This is perhaps to be expected, as estimating the first moments of the velocities is a more straightforward task than a full p.d.f., as discussed previously. All three predictions at $Re = 40$ are very close to the true DNS profiles and the mean velocity profile is predicted quite faithfully at all Re considered. Note the significant weakening of \overline{U} , which was commented on in [§ 2.2](#), consistent with the increased decoupling from the background forcing (for reference, the laminar velocity profile is proportional to Re in our scaling). However, the fluctuating velocities grow with increasing Re . Predictions for the r.m.s. velocity fluctuations suffer slightly for $Re > 40$, but errors are consistently below $\mathcal{O}(10\%)$. These velocity profile reconstructions are roughly in line with the observable-agnostic predictions at $Re \in \{40, 100\}$ by [Page *et al.* \(2024b\)](#).

Tracking how the weights for the RPOs change along their solution curves provides an indication of the changing dynamical importance of that solution as Re increases. This is shown for a sample of representative solution branches in [figure 12\(a\)](#), where D was again used as the training statistic to fit the weights. The relative weight $w_j/\max_j w_j$ for each RPO is indicated by the size of its marker (circles), while points where the weight $w_j/\max_j w_j < 10^{-4}$ are indicated with a cross. The figure indicates the same broad trends identified previously, that weights are larger when the RPO sits closer to the centre of the

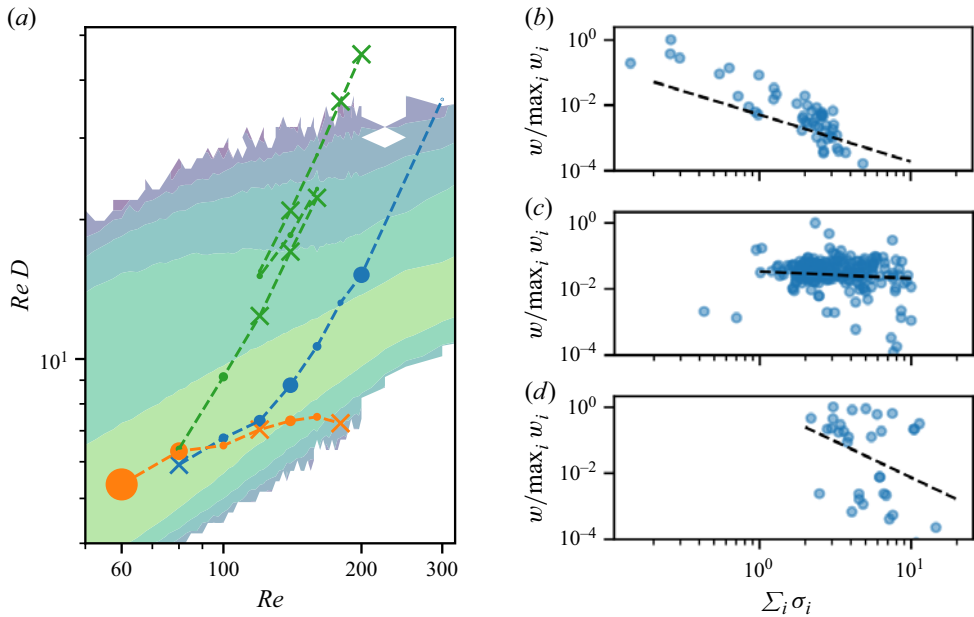


Figure 12. (a) Visualisation of the changing weights along the solution curves of some example RPOs, with D as the training statistic. The RPOs are visualised via their time-averaged, scaled dissipation rate, and the relative size of each circle along the curves indicates the relative weight $w_j / \max_j w_j$ of the RPO at that particular Re . Each colour represents a different solution branch; larger circles indicate larger weights, while crosses denote RPOs with $w_j / \max_j w_j < 10^{-4}$. (b–d) Dependence of the weights on the real part of the sum of unstable Floquet exponents $\sum_j \sigma_j$, $\sigma_j > 0$, for each RPO at $Re = 40, 100, 200$, respectively. The dotted black line denotes the line of best fit, indicating an inverse dependence of $(\sum_i \sigma_i)^{-1.43}$, $(\sum_i \sigma_i)^{-0.21}$ and $(\sum_i \sigma_i)^{-2.19}$, at $Re = 40, 100, 200$, respectively. Only the results at $Re = 40$ appear to show a clear trend. A total of 163, 22 and 28 RPOs are cut-off, respectively, due to the truncation of the weights at 10^{-4} .

turbulent distribution (see orange curve in figure 12) and drop as the RPO moves to higher dissipation values to become less dynamically relevant (see green and blue curves – these are class 1 solutions) – though note that the weights tend to vary non-monotonically along the branch, which is to be expected given the emergence/disappearance of solutions in folds and the occurrence of other bifurcations not documented/analysed here.

The data-driven approach to assigning weights to solutions is driven by an acceptance that the solution library is incomplete. One major flaw in our incomplete library of solutions is that no search was made for pre-periodic solutions with a finite number of shift-reflects in the vertical. We have also seen that many solutions which were believed to be on the attractor (overlap with turbulence in low-dimensional projections) may actually be outside, but still somehow contain ‘turbulent’ dynamical processes. An examination of the dependence of the weights on the Floquet multipliers of the RPOs (figure 12b–d) indicates the expected inverse correlation $w_j \sim (\sum_i \sigma_i)^{-1}$ (see also the discussion of Redfern *et al.* 2024) for moderate $Re = 40$, but shows little correlation at higher Re . The indication is a decoupling of the fitting from any underlying dynamical properties of the solutions, which is perhaps to be expected given both the fate of classes 1 and 3 as $Re \rightarrow \infty$ and the diminished cover of the tails of the p.d.f.s.

It is pertinent at this point to question the utility of a periodic orbit statistical reconstruction in light of these results. The library of solutions at $Re = 100$ (for example) is, to our knowledge, the largest set of (apparently ‘dynamically relevant’) solutions

assembled and is the result of thousands of GPU hours of computation (as documented by Page *et al.* 2024*a,b*) augmented by our continuation effort. The success of future efforts to accurately estimate statistics beyond mean values rests both on convergence of solutions with longer periods, bringing additional computational challenges, and an assessment of whether these states are embedded in the turbulent attractor, where we have seen that an inspection of low-dimensional projections is likely very misleading. One caveat here is that this analysis is based on two-dimensional turbulence and may not apply directly to the three-dimensional case. Given the apparent connection of many of the converged RPOs to inviscid solutions, we explore now whether the key dynamical features of the large set of $Re = 100$ solutions can be captured in a much simpler point vortex model.

4. Labelling with point vortex RPOs

Motivated by the apparent connection of the class 1 RPOs to solutions of the unforced Euler equation, and the wide range of vortex dynamics encompassed in this set of orbits, this section explores the feasibility of fitting exact point vortex solutions to the RPOs in our library. Our analysis is focused primarily on the large set of solutions converged at $Re = 100$, of which a large number achieved the class 1 $D \sim Re$ scaling on continuation, some after going through a series of fold bifurcations before moving away from ‘turbulent’ dissipation values. The motivation for doing this is the known connection of continuous families of Euler solution to point vortex dynamics (e.g. Stuart vortices – Stuart 1967; Crowdy 2003) in particular limits.

Fitting is done by matching the dynamics of the point vortices to the dynamics of the vortex cores of the turbulent RPOs, while insisting that the time evolution of the point vortices is also a relative periodic orbit. This approach relies on gradient-based optimisation of a scalar loss function which depends on entire point vortex trajectories, and which is accomplished here using a fully differentiable solver (Cleary & Page 2023).

The labelling procedure is broken down into three stages: (i) extraction of the vortex cores from the turbulent RPOs and the initialisation of a representative point vortex system (described in § 4.1); (ii) gradient-based optimisation which updates the initial positions and circulations of the point vortex system to improve the matching between the point vortex trajectories and the dynamics of the reference turbulent RPO (§ 4.2); (iii) a point vortex Newton-GMRES-Hookstep RPO solver then attempts to converge to an exact point vortex RPO (also detailed in § 4.2).

At the initialisation stage, a choice must be made for the number of point vortices N_v used in the fitting of each RPO. Representative trajectories of the doubly periodic point vortex system for various choices of N_v are shown in figure 13. The $N_v = 2$ system is integrable (Stremmer & Aref 1999): the vortices translate uniformly together – see figure 13(a). The $N_v = 3$ system is also integrable when the net circulation of the system is zero – the conserved quantities being the Hamiltonian and the horizontal and vertical momenta which commute for $\sum_j \Gamma_j = 0$ (Stremmer & Aref 1999), as is the case here to satisfy the Gauss constraint. This system was studied extensively by Stremmer & Aref (1999), who showed it can be mapped to the problem of advection of a passive particle by a system of stationary vortices, with considerably richer dynamics than the equivalent system on the unbounded plane. For example, they showed that if the ratio of the vortex circulations is rational, then the relative motion of the vortices is periodic, exhibiting so-called ‘paired’, ‘coupled’, ‘collective’ or ‘wandering’ motion. In the more general case of a non-rational ratio of vortex circulations, as shown in figure 13(b), the relative motion can be aperiodic. For higher $N_v \geq 4$, the system is not integrable and exhibits the usual features of low-dimensional Hamiltonian chaos (Aref & Pomphrey 1982).

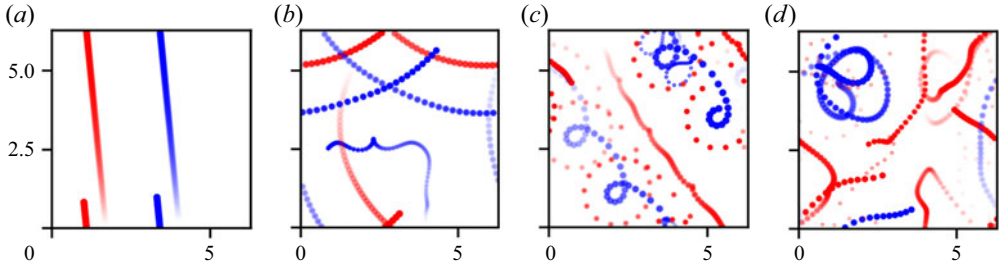


Figure 13. Sample trajectories of randomly generated point vortex systems with (a–d) $N_v = 2, 3, 4$ and 6 vortices, respectively, in a doubly periodic domain of size $2\pi \times 2\pi$, with circulations normalised such that $\max_{\alpha} |\Gamma_{\alpha}| = 10$ and $\sum_{\alpha} \Gamma_{\alpha} = 0$. Trajectories are plotted as red (blue) points denoting positive (negative) circulation, such that increasing opacity denoting increasing time and the separation of points along a trajectory gives an indication of the speed of that vortex. Each system is simulated for 30 time units and is set to have net zero circulation.

4.1. Vortex initialisation

The first stage in the labelling procedure is the initialisation of a point vortex system from the reference turbulent RPO at $Re = 100$. First, $N_S = 50$ snapshots are sampled from the reference Kolmogorov RPO, equally spaced in time. The vortex cores in each snapshot are extracted using the vortex identification method outlined in § 2.2 and the modal number of vortex cores at each snapshot throughout the period of the RPO, \bar{N}_v , is computed.

This procedure seeks to initialise a point vortex system that matches the strongest vortex cores over the period of the turbulent RPO. Trajectories of vortex cores in the reference RPO are tracked by connecting vortex cores i and j in sequential snapshots k and $k + 1$ which satisfy the similarity condition

$$\frac{\|X_i^k - X_j^{k+1}\|_2}{\|L\|_2} < 0.1, \quad (4.1)$$

where $X_i = (\tilde{x}_i, \tilde{y}_i, \Gamma_i, A_i)$ is a state vector collating the centre of vorticity, circulation and area of each vortex core. The normalisation L is set to $(L_x, L_y, \Gamma_i, A_i)$, rather than X_i . This choice of normalisation is made so that the similarity condition is independent of the location of the vortex cores on the domain, and depends instead on the inter-core distance – note that periodic boundary conditions are accounted for when computing the difference between \tilde{x}_i^k and \tilde{x}_j^{k+1} . As the turbulent solutions are typically RPOs, there is a translational discontinuity between the first and final snapshot of the period of the solution. This discontinuity is taken into account by translating the vortex cores in the first snapshot by the reverse translational shift when attempting to connect a trajectory between the first and final snapshots.

Some vortex cores will not satisfy the vorticity thresholding condition in (2.6) for all snapshots along their trajectories. In general, very few vortex-core trajectories exist over the full period of the reference RPO due to the fluctuating strength of the vortices. The trajectories of the vortex cores are sorted by their average circulation over their existence,

$$\langle \Gamma_i \rangle = \frac{\int_0^T \Gamma_i(t) \mathbb{I}_i(t) dt}{\int_0^T \mathbb{I}_i(t) dt}, \quad (4.2)$$

where the indicator function $\mathbb{I}_i(t)$ equals 1 if vortex core trajectory i exists at time t , and 0 otherwise. The average strength $\langle \Gamma_i \rangle$ is used as a measure of the dynamical importance

of the i th vortex core trajectory. The $\bar{N}_v + 1$ most dynamically important trajectories are extracted and $\langle A_i \rangle$ is also computed for each of these trajectories. The areas and circulations of the point vortex state are initialised with $\langle A_i \rangle$ and $\langle \Gamma_i \rangle$, respectively (the areas are required to compare to the reference RPO evolution – see later). If all of the extracted trajectories co-exist in some snapshot of the reference RPO, the centres of vorticity of the vortex cores in this snapshot are used to initialise the positions of the point vortices. Otherwise, the snapshot with the smallest number of absent vortices is used and the positions of the absent vortices are initialised with their average centres of vorticity $\langle \tilde{\mathbf{x}}_i \rangle$.

4.2. Convergence

The dynamical initialisation algorithm outlined in the previous subsection outputs the position, circulation and area of $\bar{N}_v + 1$ point vortices. These point vortices correspond to the $\bar{N}_v + 1$ strongest vortex core trajectories extracted from the reference RPO. However, there is no guarantee that this initialisation will be near to an exact point vortex solution. The initialisation is also very sensitive to the vorticity extraction condition in (2.6). This is because vortex cores will not satisfy (2.6) for sections of their trajectories, directly impacting the modal number of vortex cores extracted \bar{N}_v . For these reasons, we employ gradient-based optimisation to improve the point vortex fit and reduce the sensitivity to the extraction method, before converging onto an exact point vortex RPO with a Newton-GMRES-hookstep solver.

The gradient-based optimisation seeks to enforce two specific behaviours in the point vortex evolution: (i) that the point vortices must shadow the vortex cores from the reference RPO closely and (ii) that the point vortex evolution itself must also be (relative) time periodic. We construct a scalar loss function \mathcal{L} to target these two effects. Unlike in the Newton-GMRES-hookstep solver subsequently used to converge exact solutions, the circulations of the point vortices are allowed to change during optimisation.

Matching to the reference vortex cores is done by a contribution \mathcal{L}_{match} to the total scalar loss function \mathcal{L} . The contribution \mathcal{L}_{match} is included to attempt to minimise the difference between the reference RPO and a grid-based representation of the point vortex solution. The singular point vortex circulations of an N_v -vortex configuration at position $\mathbf{x}^* = [\mathbf{x}_1^*, \dots, \mathbf{x}_{N_v}^*] \in \mathbb{R}^{2N_v}$ (note $\mathbf{x}_\alpha^* := [x_\alpha^*, y_\alpha^*]$), with circulations $\boldsymbol{\Gamma} \in \mathbb{R}^{N_v}$ and areas $\mathbf{A} \in \mathbb{R}^{N_v}$, are used to define a doubly periodic observable by placing a Gaussian at each vortex location:

$$G(\mathbf{x}; \mathbf{x}^*, \boldsymbol{\Gamma}) := \sum_{\alpha=1}^{N_v} g_\alpha(\mathbf{x}; \mathbf{x}_\alpha^*, \Gamma_\alpha), \quad (4.3)$$

where

$$g_\alpha(\mathbf{x}; \mathbf{x}_\alpha^*, \Gamma_\alpha) := \sum_{n,m=-\infty}^{\infty} \tilde{g}_\alpha(\mathbf{x}; \mathbf{x}_\alpha^* + (nL, mL), \Gamma_\alpha) \quad (4.4)$$

is a doubly periodic function with Gaussian contributions from all image vortices, i.e.

$$\tilde{g}_\alpha(\mathbf{x}; \mathbf{x}_\alpha^*, \Gamma_\alpha) = \frac{\Gamma_\alpha}{2\pi\sigma_\alpha^2} \exp\left(-\frac{(\mathbf{x} - \mathbf{x}_\alpha^*)^2}{2\sigma_\alpha^2}\right). \quad (4.5)$$

In this expression, the properties of the vortex cores extracted from the reference turbulent RPO are used to define the standard deviation of the Gaussian, $\sigma_\alpha^2 = 0.1A_\alpha$. This smoothing of the vorticity delta functions centred on each point vortex allows for

comparison between snapshots of the point vortex system and the vorticity field of the reference Kolmogorov RPO. It also allows for the possibility of vortex permutations if their strengths and areas are equal. To compute the doubly periodic Gaussian distribution in (4.3), we approximate its Fourier series with the Fourier transform of the Gaussian distribution on the infinite domain

$$\frac{1}{4\pi^2} \int_{-\pi}^{\pi} \int_{-\pi}^{\pi} g_{\alpha}(\mathbf{x}) e^{-i\mathbf{k} \cdot \mathbf{x}} d^2 \mathbf{x} \approx \frac{1}{4\pi^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \tilde{g}_{\alpha}(\mathbf{x}) e^{-i\mathbf{k} \cdot \mathbf{x}} d^2 \mathbf{x}, \quad (4.6)$$

which we find to be numerically accurate to machine precision when $\sigma_{\alpha}^2 < 2$. The Fourier transform approximation $\hat{g}_{\alpha}(\mathbf{k})$ is derived by applying the shift and stretching Fourier theorems to the standard Gaussian Fourier transform

$$\hat{g}_{\alpha}(\mathbf{k}) = \frac{1}{4\pi^2} e^{-i\mathbf{k} \cdot \mathbf{x}_{\alpha}} e^{-k^2 \sigma_{\alpha}^2 / 2}. \quad (4.7)$$

Taking the inverse Fourier transform of (4.7) yields the approximation for the doubly periodic Gaussian distribution in (4.3).

The function \mathcal{L}_{match} then minimises the average L_2 distance between snapshots extracted from the turbulent RPO $\{\omega(\mathbf{x}, t) : 0 \leq t < T_{RPO}\}$ and the Gaussian observable of the point vortex solution over the period T_{RPO} of the turbulent solution,

$$\mathcal{L}_{match}(\mathbf{x}^*, \mathbf{\Gamma}) := \frac{1}{N_S} \sum_{i=1}^{N_S} \frac{\|G(\mathbf{x}; \mathbf{f}^{t_i}(\mathbf{x}^*), \mathbf{\Gamma}) - \omega(\mathbf{x}, t_i)\|^2}{\|\omega(\mathbf{x}, t_i)\|^2}, \quad (4.8)$$

where $\mathbf{f}^t(\mathbf{x}^*)$ is the time forward map of (2.15), $t_n = nT_{RPO}/N_S$ are equally spaced times over the period and the norm $\|v\|^2 := (1/V) \int_V v^2 d^2 \mathbf{x}$. The period and vortex areas A are kept fixed throughout the optimisation, so the only adjustable parameters for \mathcal{L}_{match} are the initial vortex positions \mathbf{x}^* and their circulations $\mathbf{\Gamma}$.

The second contribution to the overall loss function \mathcal{L} is designed to target point vortex RPOs, and is termed \mathcal{L}_{RPO} . To allow for permutations of the point vortices between the states \mathbf{x}^* and $\mathbf{f}^{T_{RPO}}(\mathbf{x}^*)$, this loss function is also defined in terms of the periodic Gaussian observable in (4.3),

$$\mathcal{L}_{RPO}(\mathbf{x}^*, \mathbf{\Gamma}, s) = \frac{\|G(\mathbf{x}; \mathcal{T}^s \mathbf{f}^{T_{RPO}}(\mathbf{x}^*), \mathbf{\Gamma}) - G(\mathbf{x}; \mathbf{x}^*, \mathbf{\Gamma})\|^2}{\|G(\mathbf{x}; \mathbf{x}^*, \mathbf{\Gamma})\|^2}, \quad (4.9)$$

where $\mathcal{T}^s = \mathcal{T}^{(s_x, s_y)}$ generates translational shifts in the x - and y -directions. While the turbulent RPOs cannot shift in y (we do not have solutions in our library which shift-reflect), the point vortex states are allowed to drift in both directions. The adjustable parameters are \mathbf{x}^* and $\mathbf{\Gamma}$ (as in \mathcal{L}_{match}), as well as the shifts s_x and s_y .

The full loss function is a combination of the two terms discussed previously:

$$\mathcal{L}(\mathbf{x}^*, \mathbf{\Gamma}, s) = \kappa \beta \mathcal{L}_{match}(\mathbf{x}^*, \mathbf{\Gamma}) + (1 - \kappa) \mathcal{L}_{RPO}(\mathbf{x}^*, \mathbf{\Gamma}, s), \quad (4.10)$$

where the hyperparameter $\beta = 50$ is set to a large number to increase the relative importance of the matching loss term. This large weighting of \mathcal{L}_{match} places a heavy emphasis on finding point vortex solutions which closely mimic the dominant vortex dynamics, which are themselves time periodic. The second term \mathcal{L}_{RPO} should be thought of as a small correction to aid in this convergence towards exactly recurring solutions. The hyperparameter κ is defined to conveniently set either of the loss terms to zero.

The gradients $\nabla_{\mathbf{x}^*, \mathbf{\Gamma}, s} \mathcal{L}$ are passed to an Adam optimiser (Kingma & Ba 2015) with an initial learning rate of $\eta = 10^{-2}$. The optimiser is allowed to run for a maximum

$N_{opt} = 1000$ steps, or until $\mathcal{L}_{RPO} \leq 10^{-5}$. The value of the hyperparameter κ has a strong effect on the likelihood of satisfying this tolerance on \mathcal{L}_{RPO} . For example, for $\kappa = 0.5$, just 4 % of initialisations reach this tolerance (the \mathcal{L}_{match} term is still heavily weighted due to our choice of β), compared with the 74 % of initialisations for $\kappa = 0$. All the optimiser outputs are passed to a Newton solver to attempt convergence regardless of whether they reach this tolerance level.

Two non-differentiable operations are then performed on the output of the gradient-based optimisation. The first is to remove any vortices which satisfy $|\Gamma_\alpha| < \varepsilon$, where ε is a tolerance chosen as one of $\varepsilon \in \{0, 0.01, 0.05\}$. For reference, vortex circulations were typically in the range $|\Gamma_\alpha| \in [1, 5]$ after the gradient-based optimisation. As \mathcal{L}_{RPO} is computed in terms of the Gaussian observable, the optimiser tends to reduce the circulation of any point vortices with a large contribution to \mathcal{L}_{RPO} . Removing these weak vortices improves the final convergence rate of the Newton-GMRES-hookstep solver.

The second non-differentiable operation is the identification of any vortices which permute their positions over the period of the point vortex solution. As \mathcal{L}_{RPO} is computed in terms of the grid-based Gaussian observable, vortex cores with similar areas and circulations are allowed to permute over the period T_{RPO} while maintaining small \mathcal{L}_{RPO} loss. As such, we have not yet needed to be explicit about the permutation symmetry of point vortices with equal circulation. However, the Newton-GMRES-hookstep solver used in the final convergence step finds roots of the residual defined in terms of the vortex positions

$$\mathbf{F}(\mathbf{x}^*, s, T) = \mathcal{T}^s \mathbb{P} \mathbf{f}^T(\mathbf{x}^*) - \mathbf{x}^*, \quad (4.11)$$

where the permutation \mathbb{P} must be computed *a priori*. A modified Jonker–Volgenant variant (Crouse 2016) of the Hungarian algorithm (Kuhn 1955) is used to compute the vortex permutation \mathbb{P} between \mathbf{x}^* and $\mathbf{f}^{T_{RPO}}(\mathbf{x}^*)$ after the optimisation and prior to the initialisation of the Newton solver (Cleary & Page 2023). The simplest permutation is a pairwise swap of two vortices \mathbf{x}_α^* and \mathbf{x}_β^* , which can be written as a 2-cycle, $\mathbb{P} = (\alpha \beta)$. However, higher-order k -cycles which permute k vortices are also possible in point vortex RPOs. For the state $\mathcal{T}^s \mathbb{P} \mathbf{f}^T(\mathbf{x}^*)$ to be a symmetric copy of \mathbf{x}^* , all vortices belonging to the k -cycle must share the same circulation. Any permutation can be decomposed into disjoint cycles ξ_i , such that $\mathbb{P} = \prod_i \xi_i$ and each vortex belongs to at most one ξ_i . To ensure that the permutation symmetry is respected, the circulation of all vortices belonging to each ξ_k is then manually set to their mean circulation. All vortex circulations are then fixed and not further updated by the Newton solver.

When converging solutions with the Newton solver, we apply a threshold on the Newton residual and a condition that the period is not significantly adjusted after optimisation:

$$\frac{\|\mathcal{T}^s \mathbb{P} \mathbf{f}^T(\mathbf{x}^*) - \mathbf{x}^*\|}{\|\mathbf{x}^*\|} < 10^{-8}, \quad (4.12)$$

$$\frac{|T - T_{RPO}|}{T_{RPO}} < 0.25. \quad (4.13)$$

The optimal initial size of the trust region for the Hookstep solver was found to be $0.01 \times \|\mathbf{x}^*, s, T\|_2$. This small trust region resulted in a greater rate of dynamically representative convergences – larger values resulted in point vortex RPOs that bore little resemblance to the reference solution; smaller initial trust regions required too many Newton steps for convergence.

Name	κ	ε	Total convergences	Total matches	$\mathcal{L}_{rel} < 1$
1a	1/2	0	238	117	63
1b	1/2	0.01	237	117	63
1c	1/2	0.05	238	117	63
2a	0	0.01	411	138	64
2b	0	0.05	468	146	73
3a	Anneal	0.01	255	120	67
3b	Anneal	0.05	266	122	68

Table 2. Summary of point vortex RPO-fit experiments run using 252 turbulent RPOs at $Re = 100$ to initialise the optimisation. Each reference RPO yields three different initialisations, by extracting the modal number of vortex cores, N_v , as well as $N_v \pm 1$, over the period of the reference RPO.

4.3. Optimisation results

A number of different experiments were run, varying the circulation threshold ε , controlling which vortices are retained, and the annealing parameter κ in (4.10). These experiments are summarised in table 2, while examples of some converged solutions, overlaid over the reference Kolmogorov RPO, are shown in figure 14. The labelling procedure was initialised with the 252 RPOs converged at $Re = 100$ (151 solutions from the original library, 101 new solutions from the continuation). Three different initialisations were run for each of these reference turbulent RPOs, by extracting the $N_v - 1$, N_v and $N_v + 1$ most dynamically important vortex cores over the period of the reference RPO (the overbar indicating the modal number of point vortices is dropped from now on). Vortex numbers are typically $2 \leq N_v \leq 8$, with most periodic orbits being well described with $N_v = 4$ or 5. The ‘Total convergences’ column in table 2 reports the total number of convergences (larger than 252 for each experiment due to the three initialisations run for each RPO), while the ‘Total matches’ column reports how many of the reference RPOs yielded at least one converged point vortex RPO.

Overall, table 2 indicates that all parameter settings achieve very large numbers of converged solutions, though these convergences may in some cases differ substantially from the reference simulation – point vortex states which are ‘similar’ to the underlying turbulent periodic orbit are quantified by a dynamical relevance metric $\mathcal{L}_{rel} \leq 1$, which is introduced and discussed later. In some cases in table 2, the parameter κ is labelled ‘Anneal’: we use this terminology to describe an optimisation that is initialised with $\kappa = 1$, and in which κ is then decreased by $1/N_{opt}$ at each optimisation step. This annealing aims to initially closely match the reference turbulent RPO, while slowly increasing the relative importance of the \mathcal{L}_{RPO} throughout the optimisation procedure. This scheduling results in a slight increase in the total number of point vortex RPO convergences compared with a constant equal weighting ($\kappa = 0.5$) of \mathcal{L}_{RPO} and \mathcal{L}_{match} . Ignoring the \mathcal{L}_{match} term (setting $\kappa = 1$) results in a significant increase in the total number of point vortex RPO convergences.

The solutions reported in the top three rows of figure 14 all represent cases where the converged point vortex dynamics closely track the dominant vortex cores in the original solution. They include cases featuring a multipolar structure (top row) and a bound state or triangular vortex (middle row) – in both of these examples, the point vortex RPO involves a permutation of vortices at each period. The final example is a crystal-like structure in which two rows of same-signed vortices translate to the right. In all three cases, the period of the converged point vortex solution matches the reference turbulent RPO to within a few percent or better.




Figure 14. Out-of-plane vorticity (contours) of three RPOs from Kolmogorov flow at $Re = 100$. Snapshots in each row are sampled uniformly in time along the period of that solution. Overlaid on each snapshot are the point vortices at the corresponding (proportional) time along the period of the matched point vortex RPO. The size of the markers is proportional to the circulation of the vortex, and the lines denote the path swept out by each point vortex over its full evolution. Red (white) markers denote positive (negative) circulations. The reference RPO period T_{RPO} , reference RPO shift s_{RPO} , point vortex RPO period T and point vortex shifts s for each solution are: (top row) $T_{RPO} = 1.34$, $s_{RPO} = 0.0075$, $T = 1.2$, $s = (-0.032, 0.012)$; (second row) $T_{RPO} = 1.79$, $s_{RPO} = 0.085$, $T = 1.755$, $s = (0.05, 0.006)$; (third row) $T_{RPO} = 1.27$, $s_{RPO} = 0.69$, $T = 1.27$, $s = (0.48, 0)$. This point vortex solution is a vortex crystal (relative equilibrium); (bottom row) $T_{RPO} = 4.16$, $s_{RPO} = -0.756$, $T = 4.34$, $s = (0.04, -0.12)$.

However, by visual inspection, the dynamics of some converged point vortex RPOs differs qualitatively from the reference turbulent RPO. An example of which is shown in the bottom row of figure 14, in which a six vortex system has been converged which only slightly resembles the reference RPO – there are two trapped systems of $N_v = 4$ and $N_v = 2$ vortices which roughly align with large-scale vortices in the turbulent flow.

To quantify the dynamical relevance of the point vortex state to the original turbulent dynamics, we compute the translation-and-time-reduced metric:

$$\mathcal{L}_{rel} := \min_{s, \Delta t} \frac{1}{N_S} \sum_{j=1}^{N_S} \frac{\left\| G(\mathbf{x}; \mathcal{T}^s \mathbf{f}^{t_j^v + \Delta t}(\mathbf{x}^*), \mathbf{\Gamma}) - \sum_i \mathbb{I}_i(\mathbf{x}, t_j) \omega(\mathbf{x}, t_j) \right\|^2}{\left\| \sum_i \mathbb{I}_i(\mathbf{x}, t_j) \omega(\mathbf{x}, t_j) \right\|^2}, \quad (4.14)$$




Figure 15. Histograms of the dynamical relevance metric (4.14) for experiments 1c, 2b and 3b detailed in table 2, with a bin size of 0.05. The mean and mode of each experiment are indicated by the correspondingly coloured bold dashed line and filled bars, respectively.

where the indicator function $\mathbb{I}_i(\mathbf{x}, t) = 1$ if $\mathbf{x} \in \mathcal{V}_i$ at time t – i.e. we are comparing only the vortex cores in the original RPO to the point vortex reconstruction. The solutions are compared at equally spaced times $t_j^v := jT/N_S$ and $t_j = jT_{RPO}/N_S$ along the periods of point vortex RPO T and the reference turbulent RPO T_{RPO} , respectively. Due to potential small mismatches between the alignment of the vortices in the two solutions in space and time, we search in (4.14) over a restricted range of (constant) temporal and spatial shifts: $\Delta t \in \{0, \pm 10\delta t, \pm 25\delta t, \pm 50\delta t, \pm 100\delta t\}$, and $s_x, s_y \in \{0, \pm L/100, \pm L/40, \pm L/20, \pm L/10\}$.

It is worth considering what a ‘good’ value of \mathcal{L}_{rel} would be. For unrelated point vortex observables and turbulent RPOs, we find $\mathcal{L}_{rel} \approx 2$ (the expected result given we are comparing two near-orthogonal vectors in a high-dimensional space). For the dynamically similar states obtained in figure 14, we obtain values of $\mathcal{L}_{rel} = 0.47, 0.49, 0.46$ for the first, second and third rows, respectively. The state in the bottom row of figure 14 only achieves a value of $\mathcal{L}_{rel} = 1.58$, indicating a poor representation of the turbulent dynamics. The relatively ‘large’ values of \mathcal{L}_{rel} in the well-matched configurations can be rationalised by the fact that the smoothed point vortex observable is constructed as a superposition of Gaussians, with the shape of the original vortex core modelled by a variance which is proportional to the vortex area, $\sigma_\alpha^2 = 0.1A_\alpha$: we are comparing fields which can differ quite significantly even in the best scenarios. As such, we find that $\mathcal{L}_{rel} \leq 1$ is a sensible threshold to discern whether the converged point vortex RPO captures critical features of the original turbulent solution. We also checked all point vortex states with $\mathcal{L}_{rel} \leq 1$ for uniqueness by rescaling variables such that $\max |\Gamma_j| = 1$, before comparing their rescaled periods and values of the rescaled Hamiltonian. All solutions are unique.

A comparison of the distributions of \mathcal{L}_{rel} across three representative experiments from the list in table 2 (1c, 2b and 3b) is reported in figure 15. These point vortex RPO searches all enforced removal of vortices whose circulation shrunk in the optimisation below a threshold of $\varepsilon = 0.05$, and are distinguished from one another by the relative weight placed on the two contributions to the loss (4.10): In run 1c, the ‘RPO’ contribution (4.9) was included alongside the matching constraint (4.8), while it represented the entire loss in run 2b (no matching term). In run 3b, annealing was applied to incrementally increase the contribution of the RPO term over the course of the optimisation. All of the search




Figure 16. (a) Total number of point vortex RPO convergences at each value of N_v for experiment 3b in table 2 which satisfy $\mathcal{L}_{rel} < 1$. (b) Histogram showing the distribution of \mathcal{L}_{rel} for experiment 3b as a function of N_v . The bin size for \mathcal{L}_{rel} is 0.1.

configurations summarised in figure 15 show a wide spread in \mathcal{L}_{rel} . Notably the RPO-only results (2b) show much higher \mathcal{L}_{rel} when no requirement is made that the point vortices attempt to shadow the vortex cores. We focus for the remainder of the section on the results of experiment 3b, which has the lowest mean (and mode) in \mathcal{L}_{rel} , though the results are representative of the suite of experiments in general. Note that the value of the circulation cutoff, ε , selected for analysis here (the largest, $\varepsilon = 0.05$) leads to an increased number of convergences relative to calculations which retain larger numbers of weak vortices, though the overall distribution of results are not substantially changed from those shown in figure 15.

4.4. Discussion

We now explore the properties of the library of converged point vortex states from experiment 3b and their relationship to the three classes of Kolmogorov RPO discussed in § 3.1. The number of convergences from this experiment satisfying $\mathcal{L}_{rel} < 1$ is visualised as a function of the number of point vortices N_v in figure 16(a). There is a preference for small numbers of vortices, $N_v \leq 5$, with the most solutions found for $N_v = 4$. Intriguingly, there is no clear dependence in the reconstruction error, \mathcal{L}_{rel} , on the number of vortices, N_v , other than an increase in the minimum observed value of \mathcal{L}_{rel} with increasing N_v . This is perhaps unexpected as intuitively, one might think that placing more point vortices in the system would result in a more faithful representation of the turbulent dynamics. However, the solutions in the turbulent RPO library at $Re = 100$ tend to be dominated by a small number of vortex cores (see for example figures 5 and 6), and additional point vortices are attempting to represent weak, small-scale filamentary structures which are poorly represented with Gaussians (discussed further later).

Some examples of solutions in the most common $N_v = 4$ class of point vortex RPOs are reported in figure 17. Most $N_v = 4$ solutions appear as an approximately equispaced lattice, with two dominant vortices of opposite circulation. An example of one of these ‘quasi-crystal’ configurations is shown in figure 17(a). Other representative solutions for




Figure 17. Evolution of three samples point vortex RPOs with $N_v = 4$. Trajectories are plotted as red (blue) points denoting positive (negative) circulation, such that increasing opacity denoting increasing time along the period and the separation of points along a trajectory gives an indication of the speed of that vortex. Streamlines of the final (most opaque) vortex configuration are also shown, such that darker streamlines denote greater local induced speed. Each system is simulated for one complete period. (a) Vortex crystal configuration with $T = 2.16$, $s = (-0.008, 0.002)$. (b) Tripolar configuration with $T = 1.75$, $s = (-0.056, 0.073)$. (c) Uniform crystal lattice configuration with $T = 1.27$, $s = (0.48, 0)$.

the $N_v = 4$ system are shown in the remaining two panels, and show a ‘tripole’ structure and an oscillatory lattice, featuring rows of equal-signed vortices, respectively.

A key question is whether the successful fitting of a point vortex solution to a Kolmogorov RPO says anything about the possible connection of the original RPO to a solution of the (unforced) Euler equation. To explore this point, we first examine how many of the subset of RPOs which were successfully labelled with a ‘representative’ point vortex solution at $Re = 100$ (quantified by $\mathcal{L}_{rel} < 1$) can be continued to higher Re . This is done via the histogram in figure 18(a), which compares the number of labelled branches to the total number of solutions as Re varies; both scale similarly with Re .

More revealing is figure 18(b), which identifies point-vortex-labelled RPOs on a production/dissipation plot at $Re = 100$, where the solutions are coloured according to the class assigned in § 3.1. The figure appears to be dominated by solutions from classes 1 (red, dissipation stronger than turbulent p.d.f.) and 2 (blue, dissipation within turbulent p.d.f.). However, this figure is distorted by the large number of solutions in class 2, while the labelling success rate as a percentage of the cardinality of the class was significantly larger for class 1 and class 3. These statistics are summarised in table 3, which reports the number and percentage of successfully labelled solution branches (one branch may contain multiple solutions at a given Re) for each of the three distinct branch classes. The branches which could be continued the highest in Re by the arclength continuation (class 3) were the most likely to be successfully labelled with a point vortex RPO (albeit with only a small sample size). The RPOs in class 1 were the second most likely to be successfully labelled.

Even if a solution is successfully labelled, another important consideration is whether the extracted ‘vortices’ which are used to define the point vortex initialisation (and subsequently measure the success of the ‘reconstruction’) remain as Re increases, or whether they thin and shrink. The latter behaviour is expected to occur for filamentary structures where point vortices are a poor modelling choice (see figure 1).

To probe this, we report in figure 19 p.d.f.s of the extracted vortex areas, $\{A\}$, for the Kolmogorov RPOs for which a point vortex solution was also found. P.d.f.s are shown at both $Re = 100$ where the extraction + fit was performed (colour) and at the highest Re values on the solution curve (grey). The results show that the area of the vortices for class 1 solutions is roughly maintained (and on average increases slightly)




Figure 18. (a) Histogram of all the distinct RPO branches which cross $Re = 100$ as a function of Re in grey. The histogram of the branches which were successfully matched at $Re = 100$ is shown in red. (b) Energy dissipation rate D against energy production rate I at $Re = 100$, both normalised by the laminar dissipation $D_{lam} = Re/(2n^2)$, for the turbulent RPOs which were successfully matched with a point vortex RPO. The RPOs are coloured according to their branch class: 1, red; 2, blue and 3, green. The grey background is the p.d.f. computed from a trajectory of 10^5 samples, separated by 1 advective time unit. The contour levels of the p.d.f. are spaced logarithmically.

	$Re = 100$		$Re > 200$	
RPO class	Branches labelled	Percent labelled	Branches labelled	Percent labelled
1	21	33.3 %	14	43.75 %
2	32	18.4 %	4	25 %
3	7	46.7 %	6	50 %

Table 3. Number and percentage of successfully labelled ($\mathcal{L}_{rel} < 1$) branches for each of the three distinct classes of solution branches identified in § 3.1. The $Re = 100$ columns report these figures for experiment 3b, run on the library of 252 RPOs at $Re = 100$. The $Re > 200$ columns report these figures for the same labelling procedure, but the reference solutions considered are instead all the terminal branch solutions which were continued to above $Re = 200$.

on continuation. However, the class 2/3 solutions exhibit on average decreasing vortex area – though the continuation does not generally proceed particularly high in Re for most class 2 states. Class 1 also indicates a bimodal distribution with a peak at a large vortex area.

The results in figure 19 should be interpreted alongside the labelled vortex areas shown earlier in figure 1: filamentary structures at $Re = 100$ tend to have $A \lesssim 0.2$ with some exceptions, while these structures are not strong enough to be flagged at the higher value of $Re = 400$. The behaviour of the class 1 area p.d.f.s under increasing Re is consistent with an absence of filaments and the dominance of larger-area vortex cores, a result supported by two examples from both classes 1 and 3 shown in figure 20, where the point vortex solution is overlaid. Also included is a higher- Re state from the same solution branch. In the cases from class 1, the point vortex periodic orbits correspond to features




Figure 19. P.d.f.s of vortex areas (thresholding used to define this variable is described in § 2.2) in the converged RPOs for which a point vortex RPO was found. (a) Class 1 solutions: red is the p.d.f. at the starting value $Re = 100$, grey is the p.d.f. obtained using solutions at the highest Re on the branch. (b) Classes 2 and 3: cyan is the p.d.f. at the starting value of $Re = 100$, grey the p.d.f. at the highest Re on the branch (note only states which were continued successfully beyond $Re > 150$ were included in the class 2/3 results here).




Figure 20. (a–d) Contours of vorticity for four RPOs in Kolmogorov flow. Top row shows a snapshot at $Re = 100$ along with the point vortex RPO overlaid as in previous figures. The bottom row is the final state obtained after arclength continuation up in Re . The two leftmost RPOs belong to ‘class 1’, while the two rightmost RPOs belong to ‘class 3’ as described in § 3.1. The reference RPO period and shift, T_{RPO} , s_{RPO} , point vortex RPO period T and shifts s for each labelled solution on the top row are: (a) $T_{RPO} = 2.99$, $s_{RPO} = -0.44$, $T = 2.97$, $s = (-0.041, -0.023)$; (b) $T_{RPO} = 1.97$, $s_{RPO} = -0.43$, $T = 1.89$, $s = (0.01, -0.01)$; (c) $T_{RPO} = 4.61$, $s_{RPO} = 0.071$, $T = 4.695$, $s = (-0.008, -0.001)$; (d) $T_{RPO} = 4.67$, $s_{RPO} = -0.17$, $T = 4.497$, $s = (0.016, -0.018)$. The point vortex solution in panel (d) is a travelling wave. The final solutions on the bottom row were converged at $Re = 185.69$, 175.003 , 345.297 and 277.23 , from panel (a–d), respectively.

of the turbulent solution which remain as Re grows while the thin background filaments weaken in comparison. For instance, note the multipole structure in the second column approximating ring of vorticity, which amplifies under increasing Re , while the weaker triangular vortex in the same snapshot is not represented in the point vortex solution,

and weakens as Re increases. Similar effects are also apparent in the earlier examples in [figure 14](#) from class 1 (top 3 rows of that figure).

In contrast, the solutions in class 3 have resulted in the initialisation of weaker point vortices on small filamentary regions (e.g. see the weaker pair of vortices in the second class 3 example in [figure 20\(d\)](#)), a feature hinted at by the relative importance of low area vortices in [figure 19](#). In these cases, it appears we have found a point vortex RPO which can mimic the weaker viscous dynamics, without implying anything stronger about the behaviour as $Re \rightarrow \infty$.

The identification by the point vortex fit of structures in class 1 which persist as Re increases is consistent with the connection of these RPOs to exact solutions of the Euler equation as $Re \rightarrow \infty$, as suggested by their asymptotic scalings. Under continuation, these RPOs tend to be increasingly dominated by large-scale, coherent vortices while weaker filamentary structures lessen. The vortex extraction algorithm largely avoids the filamentary components for class 1 states and the point vortex fit captures the large-scale flow features which are retained in the high- Re limit. Interestingly, these class 1 solutions contain examples with many interacting vortices (typical of high dissipation states seen in past work (Chandler & Kerswell 2013; Page *et al.* 2024a; Redfern *et al.* 2024)) but also examples of condensate-like structures dominated by a pair of opposite signed vortices, which are reminiscent of the Euler solutions discussed in previous work exploring Euler/Navier–Stokes connections (Kim & Okamoto 2015; Kim *et al.* 2017; Zhigunov & Grigoriev 2023). The weaker performance of the point vortex fit in classes 2 and 3, and failure to extract the structure retained as Re increases, reflects the fact that – particularly in class 2 – the flow field is dominated by a large scale vortex pair and thin filamentary structures.

5. Conclusion

This study presented a continuation effort of a large number of numerically exact recurrent solutions in two-dimensional Kolmogorov flow, over the range $Re \in (30, 1000)$. Libraries of Kolmogorov RPOs at different Re were constructed from sets of ‘starting’ solutions at $Re = 40$ and $Re = 100$, and used to reconstruct summary statistics of the turbulence. The reconstruction was performed with a data-driven approach in which the weights in an RPO-expansion were determined from a fit to a single p.d.f. The quality of the reconstruction deteriorated dramatically with changing Re , while the numerically assigned weights become poorly correlated with the instability of the underlying RPOs. This raises serious questions about the use of RPOs to estimate statistics – particularly full p.d.f.s – given the computational expense required to assemble the starting solution libraries used here (thousands of GPU hours, see Page *et al.* 2024a,b). Further work should explore if there are more ‘robust’, perhaps longer-period, solutions with high dissipation excursions which can be used to say something about bursting events via continuation. It may also be worth considering alternative strategies for assigning weights to the periodic orbits too. For instance, the recent work of Pughe-Sanford *et al.* (2025) has described a simple method for statistical reconstruction with only a handful of periodic solutions of the Lorenz equations. Whether these ideas, or others like them, can yield robust statistical reconstruction in turbulence will determine the utility of the continued hunt for new solutions. Finally, it remains to be seen whether the conclusions drawn here are also directly relevant to three-dimensional, wall-bounded flows.

The dependence of dissipation rate on Re was used to delineate the large number of solution branches into three classes: class 1 RPOs were characterised by a dissipation rate scaling $D \sim Re$, and an analysis of the scaling of flow variables indicated a direct

connection to solutions of the unforced Euler equation as $Re \rightarrow \infty$. The RPOs on these branches feature strong vortical structures that appeared largely decoupled from the specifics of the Kolmogorov forcing and generalise the Euler connection to states beyond the ‘unimodal’ structures considered previously (Kim & Okamoto 2015; Kim *et al.* 2017). Two other classes of solutions were identified which scaled either with the average turbulent dissipation (class 2, $D \sim Re^{-1/2}$) or weaker than it (class 3, $D \sim Re^{-1}$). Class 2 solutions remained contained within a finite- Re region, many repeatedly folding back and forth under continuation. Class 3 solutions show a scaling consistent with a connection to a forced Euler equation in the high- Re limit. Both class 1 and 3 solutions are particularly remarkable: their appearance in low-dimensional projections at the Re values where they were first obtained seems to suggest that many are ‘dynamically relevant’ to the turbulence. It is only under continuation that their apparent irrelevance (they may not be on the chaotic attractor) is revealed.

Finally, motivated by the connection of large numbers of RPOs to inviscid solutions, a method was introduced to search for point vortex relative periodic orbits that reproduced the dominant vortex interactions in the original solution library at $Re = 100$. This was done via gradient-based minimisation of a scalar loss function. The method proved successful for a wide range of RPOs, but only for solutions in class 1 did it identify vortical structures that persist under increasing Re . The success can perhaps be attributed to the fact that these states become dominated by (smooth) large-scale structures as Re increases at the expense of small scale filamentary vortices. The success of the fit is consistent with a connection to exact inviscid solutions.

Acknowledgements. We are grateful to the referees for their detailed comments which have helped substantially to clarify the messages in this paper. This work has made use of the resources provided by the Edinburgh Compute and Data Facility (ECDF) (<https://www.ecdf.ed.ac.uk/>).

Funding. This research has been supported by the UK Engineering and Physical Sciences Research Council through the MAC-MIGS Centre for Doctoral Training (EP/S023291/1). J.P. acknowledges support from UKRI Frontier Guarantee Grant EP/Y004094/1.

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

REFERENCES

- AREF, H. & BUREN, M. 2005 Vortex Triple Rings. *Phys. Fluids* **17**, 057104.
- AREF, H., NEWTON, P.K., STREMLER, M.A., TOKIEDA, T. & VAINCHTEIN, D.L. 2003 Vortex crystals. In *Advances in Applied Mechanics*, pp. 1–79. Elsevier.
- AREF, H. & POMPHREY, N. 1982 Integrable and chaotic motions of four vortices I. The case of identical vortices. *Proc. R. Soc. Lond. A Math. Phys. Sci.* **380** (1779), 359–387.
- ARTUSO, R., AURELL, E. & CVITANOVIC, P. 1990a Recycling of strange sets: I. Cycle expansions. *Nonlinearity* **3** (2), 325.
- ARTUSO, R., AURELL, E. & CVITANOVIC, P. 1990b Recycling of strange sets: II. Applications. *Nonlinearity* **3** (2), 361.
- BATCHELOR, G.K. 1967 *An Introduction to Fluid Dynamics*. Cambridge University Press.
- BATCHELOR, G.K. 1956 On steady laminar flow with closed streamlines at large Reynolds number. *J. Fluid Mech.* **1** (2), 177–190.
- BENZI, R., COLELLA, M., BRISCOLINI, M. & SANTANGELO, P. 1992 A simple point vortex model for two-dimensional decaying turbulence. *Phys. Fluids A: Fluid Dyn.* **4** (5), 1036–1039.
- BENZI, R., PATARNELLO, S. & SANTANGELO, P. 1987 On the statistical properties of two-dimensional decaying turbulence. *Eur. Phys. Lett.* **3** (7), 811.
- BRADBURY, J. *et al.* 2018 JAX: composable transformations of Python+NumPy programs, 0.3.13. Available at: <http://github.com/google/jax>.
- BUDANUR, N.B., BORRERO-ECHEVERRY, D. & CVITANOVIĆ, P. 2015 Periodic orbit analysis of a system with continuous symmetry—a tutorial. *Chaos* **25** (7), 073112.

- BUDANUR, N.B., SHORT, K.Y., FARAZMAND, M., WILLIS, A.P. & CVITANOVIĆ, P. 2017 Relative periodic orbits form the backbone of turbulent pipe flow. *J. Fluid Mech.* **833**, 274–301.
- CAMPBELL, L.J. & ZIFF, R.M. 1979 Vortex patterns and energies in a rotating superfluid. *Phys. Rev. B* **20** (5), 1886.
- CARNEVALE, G.F. & KLOOSTERZIEL, R.C. 1994 Emergence and evolution of triangular vortices. *J. Fluid Mech.* **259**, 305–331.
- CARTON, X.J., FLIERL, G.R. & POLVANI, L.M. 1989 The generation of tripoles from unstable axisymmetric isolated vortex structures. *Eur. Phys. Lett.* **9** (4), 339.
- CHANDLER, G.J. & KERSWELL, R.R. 2013 Invariant recurrent solutions embedded in a turbulent two-dimensional Kolmogorov flow. *J. Fluid Mech.* **722**, 554–595.
- CHRISTIANSEN, F., CVITANOVIĆ, P. & PUTKARADZE, V. 1997 Spatiotemporal chaos in terms of unstable recurrent patterns. *Nonlinearity* **10** (1), 55.
- CLEARY, A. & PAGE, J. 2023 Exploring the free-energy landscape of a rotating superfluid. *Chaos* **33** (10), 103123.
- CROUSE, D.F. 2016 On implementing 2D rectangular assignment algorithms. *IEEE Trans. Aerosp. Electron. Syst.* **52** (4), 1679–1696.
- CROWDY, D. 1999 A class of exact multipolar vortices. *Phys. Fluids* **11** (9), 2556–2564.
- CROWDY, D.G. 2003 Stuart vortices on a sphere. *J. Fluid Mech.* **498**, 381–402.
- CVITANOVIĆ, P., ARTUSO, R., MAINIERI, R., TANNER, G. & VATTAY, G. 2016 *Chaos: Classical and Quantum*. Niels Bohr Inst.
- CVITANOVIĆ, P. 1991 Periodic orbits as the skeleton of classical and quantum chaos. *Physica D: Nonlinear Phenom.* **51** (1), 138–151.
- CVITANOVIĆ, P. 2013 Recurrent flows: the clockwork behind turbulence. *J. Fluid Mech.* **726**, 1–4.
- DEGUCHI, K. 2015 Self-sustained states at Kolmogorov microscale. *J. Fluid Mech.* **781**, R6.
- DRESCHNER, G., KOCHKOV, D., NORGAARD, P., ZEPEDA-NÚÑEZ, L., SMITH, J.A., BRENNER, M.P. & HOYER, S. 2022 Learning to correct spectral methods for simulating turbulent flows. [arXiv:2207.00556](https://arxiv.org/abs/2207.00556).
- ECKHARDT, B. & ZAMMERT, S. 2018 Small scale exact coherent structures at large Reynolds numbers in plane Couette flow. *Nonlinearity* **31** (2), R66–R77.
- FAISST, H. & ECKHARDT, B. 2003 Traveling waves in pipe flow. *Phys. Rev. Lett.* **91**, 224502.
- FARAZMAND, M. 2016 An adjoint-based approach for finding invariant solutions of Navier–Stokes equations. *J. Fluid Mech.* **795**, 278–312.
- GALLET, B. & YOUNG, W.R. 2013 A two-dimensional vortex condensate at high Reynolds number. *J. Fluid Mech.* **715**, 359–388.
- GIBSON, J.F., HALCROW, J. & CVITANOVIĆ, P. 2008 Visualizing the geometry of state space in plane Couette flow. *J. Fluid Mech.* **611**, 107–130.
- GRAHAM, M.D. & FLORYAN, D. 2021 Exact coherent states and the nonlinear dynamics of wall-bounded turbulent flows. *Annu. Rev. Fluid Mech.* **53** (1), 227–253.
- HALCROW, J., GIBSON, J.F., CVITANOVIĆ, P. & VISWANATH, D. 2009 Heteroclinic connections in plane Couette flow. *J. Fluid Mech.* **621**, 365–376.
- HALL, P. & SHERWIN, S. 2010 Streamwise vortices in shear flows: harbingers of transition and the skeleton of coherent structures. *J. Fluid Mech.* **661**, 178–205.
- HALL, P. & SMITH, F.T. 1991 On strongly nonlinear vortex/wave interactions in boundary-layer transition. *J. Fluid Mech.* **227**, 641–666.
- HOF, B., VAN DOORNE, C.W.H., WESTERWEEL, J., NIEUWSTADT, F.T.M., FAISST, H., ECKHARDT, B., WEDIN, H., KERSWELL, R.R. & WALEFFE, F. 2004 Experimental observation of nonlinear traveling waves in turbulent pipe flow. *Science* **305** (5690), 1594–1598.
- JIMÉNEZ, J. 2020 Dipoles and streams in two-dimensional turbulence. *J. Fluid Mech.* **904**, A39.
- KAWAHARA, G. & KIDA, S. 2001 Periodic motion embedded in plane Couette turbulence: regeneration cycle and burst. *J. Fluid Mech.* **449**, 291–300.
- KAWAHARA, G., UHLMANN, M. & VAN VEEN, L. 2012 The significance of simple invariant solutions in turbulent flows. *Annu. Rev. Fluid Mech.* **44** (1), 203–225.
- KERSWELL, R.R. 2005 Recent progress in understanding the transition to turbulence in a pipe. *Nonlinearity* **18** (6), R17.
- KIM, S.-C., MIYAJI, T. & OKAMOTO, H. 2017 Unimodal patterns appearing in the two-dimensional Navier–Stokes flows under general forcing at large Reynolds numbers. *Eur. J. Mech. (B/Fluids)* **65**, 234–246.
- KIM, S.-C. & OKAMOTO, H. 2010 Vortices of large scale appearing in the 2D stationary Navier–Stokes equations at large Reynolds numbers. *Japan J. Ind. Appl. Maths* **27**, 47–71.
- KIM, S.-C. & OKAMOTO, H. 2015 Unimodal patterns appearing in the Kolmogorov flows at large Reynolds numbers. *Nonlinearity* **28** (9), 3219.

- KINGMA, D.P. & BA, J. 2014 Adam: A method for stochastic optimization. [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- 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.
- KRYGIER, M.C., PUGHE-SANFORD, J.L. & GRIGORIEV, R.O. 2021 Exact coherent structures and shadowing in turbulent Taylor–Couette flow. *J. Fluid Mech.* **923**, A7.
- KUHN, H.W. 1955 The Hungarian method for the assignment problem. *Nav. Res. Logistics Q.* **2** (1–2), 83–97.
- LEWIS, D. & RATIU, T. 1996 Rotatingn-gon/kn-gon vortex configurations. *J. Nonlinear Sci.* **6** (5), 385–414.
- LUCAS, D. & KERSWELL, R. 2017 Sustaining processes from recurrent flows in body-forced turbulence. *J. Fluid Mech.* **817**, R3.
- LUCAS, D. & KERSWELL, R.R. 2015 Recurrent flow analysis in spatiotemporally chaotic 2-dimensional Kolmogorov flow. *Phys. Fluids* **27** (4), 045106.
- MCCORMACK, M., CAVALIERI, A.V.G. & HWANG, Y. 2024 Multi-scale invariant solutions in plane Couette flow: a reduced-order model approach. *J. Fluid Mech.* **983**, A33.
- MONTEMURO, B., WHITE, C.M., KLEWICKI, J.C. & CHINI, G.P. 2020 A self-sustaining process theory for uniform momentum zones and internal shear layers in high Reynolds number shear flows. *J. Fluid Mech.* **901**, A28.
- MOREL, Y.G. & CARTON, X.J. 1994 Multipolar vortices in two-dimensional incompressible flows. *J. Fluid Mech.* **267**, 23–51.
- OKAMOTO, H. 1994 A variational problem arising in the two-dimensional Navier–Stokes equations with vanishing viscosity. *Appl. Maths Lett.* **7** (1), 29–33.
- PAGE, J., BRENNER, M.P. & KERSWELL, R.R. 2021 Revealing the state space of turbulence using machine learning. *Phys. Rev. Fluids* **6**, 034402.
- PAGE, J., HOLEY, J., BRENNER, M.P. & KERSWELL, R.R. 2024a Exact coherent structures in two-dimensional turbulence identified with convolutional autoencoders. *J. Fluid Mech.* **991**, A10.
- PAGE, J., NORGAARD, P., BRENNER, M.P. & KERSWELL, R.R. 2024b Recurrent flow patterns as a basis for two-dimensional turbulence: predicting statistics from structures. *Proc. Natl Acad. Sci. USA* **121** (23), e2320007121.
- PARK, J.S. & GRAHAM, M.D. 2015 Exact coherent states and connections to turbulent dynamics in minimal channel flow. *J. Fluid Mech.* **782**, 430–454.
- PUGHE-SANFORD, J.L., QUINN, S., BALABANSKI, T. & GRIGORIEV, R.O. 2025 Computing chaotic time-averages from few periodic or non-periodic orbits. *Chaos* **35** (6), 063133.
- REDFERN, E.M., LAZER, A.L. & LUCAS, D. 2024 Dynamically relevant recurrent flows obtained via a nonlinear recurrence function from two-dimensional turbulence. *Phys. Rev. Fluids* **9**, 124401.
- REYNOSO, M., ZHIGUNOV, D. & GRIGORIEV, R.O. 2024 Self-similarity and the direct (enstrophy) cascade in forced two-dimensional fluid turbulence. *J. Fluid Mech.* **995**, A4.
- SAKAJO, T. 2019 Exact solution to a Liouville equation with Stuart vortex distribution on the surface of a torus. *Proc. Math. Phys. Engng Sci.* **475** (2224), 1–16.
- SMITH, L.M. & YAKHOT, V. 1993 Bose condensation and small-scale structure generation in a random force driven 2D turbulence. *Phys. Rev. Lett.* **71**, 352–355.
- STIELTJES, T.J. 1900 Sur certains polynômes: Qui vérifient une équation différentielle linéaire du second ordre et sur la theorie des fonctions de Lamé. *Acta Mathematica* **6**, 321–326.
- STREMLER, M.A. & AREF, H. 1999 Motion of three point vortices in a periodic parallelogram. *J. Fluid Mech.* **392**, 101–128.
- STUART, J.T. 1967 On finite amplitude oscillations in laminar mixing layers. *J. Fluid Mech.* **29** (3), 417–440.
- SURI, B., KAGEORGE, L., GRIGORIEV, R.O. & SCHATZ, M.F. 2020 Capturing turbulent dynamics and statistics in experiments with unstable periodic orbits. *Phys. Rev. Lett.* **125**, 064501.
- SURI, B., TITHOF, J., GRIGORIEV, R.O. & SCHATZ, M.F. 2017 Forecasting fluid flows using the geometry of turbulence. *Phys. Rev. Lett.* **118** (11), 114501.
- SURI, B., TITHOF, J., GRIGORIEV, R.O. & SCHATZ, M.F. 2018 Unstable equilibria and invariant manifolds in quasi-two-dimensional Kolmogorov-like flow. *Phys. Rev. E* **98** (2), 023105.
- WALEFFE, F. 2001 Exact coherent structures in channel flow. *J. Fluid Mech.* **435**, 93–102.
- WALEFFE, F. 2003 Homotopy of exact coherent structures in plane shear flows. *Phys. Fluids* **15** (6), 1517–1534.
- WANG, J., GIBSON, J. & WALEFFE, F. 2007 Lower branch coherent states in shear flows: transition and control. *Phys. Rev. Lett.* **98**, 204501.
- WEDIN, H. & KERSWELL, R.R. 2004 Exact coherent structures in pipe flow: travelling wave solutions. *J. Fluid Mech.* **508**, 333–371.

- WEISS, J.B. & MCWILLIAMS, J.C. 1991 Nonergodicity of point vortices. *Phys. Fluids A: Fluid Dyn.* **3** (5), 835–844.
- YALNIZ, G., HOF, B. & BUDANUR, N.B. 2021 Coarse graining the state space of a turbulent flow using periodic orbits. *Phys. Rev. Lett.* **126**, 244502.
- YANG, Q., WILLIS, A.P. & HWANG, Y. 2019 Exact coherent states of attached eddies in channel flow. *J. Fluid Mech.* **862**, 1029–1059.
- ZHIGUNOV, D. & GRIGORIEV, R.O. 2023 Exact coherent structures in fully developed two-dimensional turbulence. *J. Fluid Mech.* **970**, A18.