

Joint reduced model for the laminar and chaotic attractors in plane Couette flow

Bálint Kaszás¹  and George Haller¹ 

¹Institute for Mechanical Systems, ETH Zürich Leonhardstrasse 21, Zürich 8092, Switzerland

Corresponding author: Bálint Kaszás, bkaszas@stanford.edu

(Received 15 May 2025; revised 30 September 2025; accepted 30 September 2025)

We use the theory of spectral submanifolds (SSMs) to develop a low-dimensional reduced-order model for plane Couette flow restricted to the shift–reflect invariant subspace in the permanently chaotic regime at $Re = 187.8$ studied by Kreilos & Eckhardt (2012, *Chaos: Interdisciplinary J. Nonlinear Sci.*, vol. 22, 047505). Our three-dimensional model is obtained by restricting the dynamics to the slowest mixed-mode SSM of the edge state. We show that this results in a nonlinear model that accurately reconstructs individual trajectories, representing the entire chaotic attractor and the laminar dynamics simultaneously. In addition, we derive a two-dimensional Poincaré map that enables the rapid computation of the periodic orbits embedded in the chaotic attractor.

Key words: transition to turbulence, low-dimensional models, chaos

1. Introduction

Certain shear flows, such as pipe flow (Avila, Barkley & Hof 2023) and plane Couette flow (Eckhardt 2018), exhibit subcritical transition to turbulence, which results in a coexistence of an extended turbulent state with a stable laminar state. The boundary between the turbulent and laminar behaviours is often called the edge of chaos (Skufca, Yorke & Eckhardt 2006). For Couette flow in small periodic domains, this edge of chaos is the stable manifold of an unstable exact coherent state (ECS) (Waleffe 2001), called the edge state (Wang, Gibson & Waleffe 2007).

It is known that ECSs play an important role in the turbulent dynamics. The simplest solutions are fixed points that were calculated by Nagata (1990) and Gibson, Halcrow & Cvitanović (2009) in plane Couette flow. Travelling waves were discovered in pipe flow by Faisst & Eckhardt (2003) and Wedin & Kerswell (2004), and in plane Couette flow by Gibson *et al.* (2009). Periodic orbits and relative periodic orbits found by Viswanath (2007), Willis, Cvitanović & Avila (2013) and others have been demonstrated to be

embedded in the turbulent state (Budanur *et al.* 2017). Furthermore, complex behaviour tends to develop through bifurcations of simpler ECSs. In particular, a period doubling cascade (Moore *et al.* 1983), or the bifurcation of invariant tori (Ruelle & Takens 1971; Gollub & Swinney 1975) or intermittent behaviour (Pomeau & Manneville 1980) precedes the appearance of chaos (Eckmann 1981).

For these reasons, great effort has been directed towards identifying unstable ECSs in direct numerical simulations (DNS) (Kawahara & Kida 2001; Graham & Floryan 2021; Page *et al.* 2024) and in experiments (Suri *et al.* 2019).

Developing predictive dynamical models is even more challenging. A frequently used, equation-driven method of model reduction employs Galerkin projection to a low-dimensional linear subspace. However, the optimal selection of modes spanning that subspace is often non-trivial, hence a large number of modes might be necessary to represent the turbulent statistics with sufficient accuracy (Cavaliere & Nogueira 2022). In addition, correspondence between trajectories of the full and reduced models is not guaranteed, since the low-dimensional subspace is generally not invariant.

Among data-driven approaches, linear methods such as dynamic mode decomposition (Schmid 2010) are necessarily inapplicable due to the nonlinearisable nature of turbulent flows, as shown by Page & Kerswell (2019). Recently, advances in machine learning have yielded a promising alternative. Specifically, the deep-learning-based DManD method (Linot & Graham 2020) captures turbulent statistics and can also represent ECSs not enforced in their training (Linot & Graham 2023). However, deep learning methods generally suffer from a need for large amounts of training data, a time-consuming training process, and an *a priori* unclear choice of hyperparameters.

In addition to neural ordinary differential equation (ODE) based methods such as DManD, alternative machine learning approaches have been developed to model the dynamics in a latent space. A notable example is reservoir computing, and echo state networks (ESNs), in particular. These networks are universal approximators (Jaeger 2001; Ahmed, Tennie & Magri 2025), and the associated optimisation problem is quadratic and requires no iterative training. Reservoir computing has shown success in modelling spatio-temporal chaotic systems (Pathak *et al.* 2017, 2018). Recently, it has also been coupled with dimension reduction by Racca, Doan & Magri (2023), who introduced the convolutional autoencoder-echo state network (CAE-ESN). Their approach has successfully modelled Kolmogorov flow, the flow in a minimal channel, and other chaotic partial differential equations (PDEs), accurately reproducing statistical quantities and Lyapunov spectra (Margazoglou & Magri 2023; Özalp & Magri 2025).

In contrast, reducing a shear flow to a low-dimensional, attracting and structurally stable invariant manifold in its phase space offers a mathematically exact and robust construction of a reduced-order model. The recently introduced theory of spectral submanifolds (SSMs) (Haller 2025) targets low-dimensional attracting invariant manifolds emanating from stationary states, such as ECSs. Specifically, primary SSMs are defined as the smoothest invariant manifolds tangent to a selected spectral subspace of the linearised dynamics at the stationary state. Under mild non-resonance conditions, the existence and uniqueness of primary SSMs for attracting fixed points can be established (Haller 2025).

More recently, Haller *et al.* (2023) pointed out the existence of entire families of (secondary) SSMs having a lower degree of smoothness (fractional SSMs) or perturbing from mixed-mode spectral subspaces containing both unstable and stable modes (mixed-mode SSMs). Mixed-mode SSMs, in particular, can represent unstable but recurrent dynamics, and have been shown to support chaotic dynamics (Liu, Axås & Haller 2024; Xu *et al.* 2024), and to aid the parametrisation of the basin boundary (Kaszás & Haller 2024).

Apart from providing verifiable mathematical conditions for its applicability, SSM-based model reduction has the advantage over usual manifold learning approaches of providing the dimension of the underlying SSM *a priori*. This is in contrast to e.g. autoencoder-based methods that require careful optimisation of the latent dimension as a hyperparameter.

These features enable SSM-based model reduction to capture even chaotic attractors from data, as demonstrated by Liu *et al.* (2024). However, no data-driven reduced-order models have been constructed for fluid flows with coexisting non-trivial attractors. In this paper, we fill this gap in reduced flow modelling by constructing the slowest three-dimensional (3-D) mixed-mode SSM of the edge state in the plane Couette flow studied by Kreilos & Eckhardt (2012).

We show that this SSM-reduced model captures both the chaotic attractor and the laminar state of the flow with the SSM acting as an inertial manifold (Foias, Sell & Temam 1988; Liu *et al.* 2024). Our approach, therefore, extends the results of Kaszás *et al.* (2022), who obtained an SSM-reduced model for plane Couette flow with only periodic attractors, and also justifies the empirical low-dimensional model of Kreilos & Eckhardt (2012).

2. Set-up

We focus on the chaotic dynamics observed in the plane Couette flow, which is the incompressible flow in a channel between two infinite plates moving in opposite directions with velocity $\pm U$. The channel is defined as the domain $\Omega = \{(x, y, z) \in \mathbb{R}^3 : [0, L_x] \times [-h, h] \times [0, L_z]\}$, in which the velocity field $\mathbf{u} = [u, v, w](x, y, z, t)$ and the pressure p satisfy the Navier–Stokes equations

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \nu \Delta \mathbf{u}, \quad \nabla \cdot \mathbf{u} = 0, \quad (2.1)$$

where ν is the kinematic viscosity. The main parameter of the problem is the Reynolds number, defined as $Re = Uh/\nu$, where $2h$ is the distance between the moving walls. This also sets the relevant time unit as h/U .

We work with a streamwise- and spanwise-periodic computational domain corresponding to the minimal flow unit studied by Kreilos & Eckhardt (2012) and Kreilos, Eckhardt & Schneider (2014), and fix $h = 1$, $L_x = 2\pi$ and $L_z = \pi$. This domain is also comparable to those used by Nagata (1990), Page & Kerswell (2019), Linot & Graham (2023) and Kaszás *et al.* (2022).

We simulate the flow using the open source `Channelflow` library (Gibson *et al.* 2019), which employs a pseudo-spectral discretisation using Fourier modes in the streamwise and spanwise directions, and Chebyshev modes in the wall-normal direction. For a direct comparison with earlier studies, we fix the number of streamwise and spanwise modes as $N_x = N_z = 32$, and the number of wall-normal modes as $N_y = 33$. The boundary and incompressibility conditions are enforced using the influence-matrix method and tau correction (Kleiser & Schumann 1980). The time stepping algorithm implemented in `Channelflow` uses semi-implicit backward differentiation (Gibson 2014; Gibson *et al.* 2019), and generates the forward time flow map of the discretised PDE (2.1). This time evolution can be implicitly interpreted as a trajectory of a finite-dimensional dynamical system (Gibson, Halcrow & Cvitanović 2008)

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}), \quad \mathbf{x} \in \mathbb{R}^N, \quad \mathbf{f} \in \mathcal{C}^\infty, \quad (2.2)$$

with $\mathbf{x}$ denoting the collection of velocity values in physical space returned by `Channelflow`, and the total number of degrees of freedom is $N \approx O(10^5)$. The laminar

flow, which is a stable fixed point of (2.2), is expressed as $u = y$, $v = w = 0$. In the following, we focus on the dynamical system (2.2).

As shown by Kreilos & Eckhardt (2012) and Kreilos *et al.* (2014), in this computational cell, the Nagata upper and lower branch fixed points (Nagata 1990; Gibson *et al.* 2009) of (2.2) appear in a saddle-node bifurcation at $Re = 163.8$. For higher Reynolds numbers, the upper branch fixed point (UB) undergoes a Hopf bifurcation, followed by a period doubling cascade, eventually leading to chaotic dynamics.

The lower branch fixed point remains an edge state for a wide range of Reynolds numbers, and its codimension-one stable manifold forms the edge of chaos (Wang *et al.* 2007; Schneider *et al.* 2008). The chaotic attractor, which can be traced back to the upper branch fixed point, undergoes a boundary crisis (Grebogi, Ott & Yorke 1983) and disappears at $Re = 188.51$. After this point, the transient chaotic behaviour associated with turbulence is generated by a chaotic saddle (Lai & Tél 2011) and hence has a finite lifetime (Kreilos *et al.* 2014).

We focus on a Reynolds number value before the boundary crisis, $Re = 187.8$, where a genuine chaotic attractor coexists with the stable laminar state. Furthermore, we also restrict the flow to the symmetry invariant subspace of the Nagata equilibria by solving (2.1) with the shift–reflect symmetries imposed.

2.1. Spectral submanifolds

We aim to construct a simple reduced model that faithfully represents the bistability of (2.2). To this end, we seek a low-dimensional invariant manifold containing the edge state, labelled $\mathbf{x}_{LB}$. The linearised dynamics around the edge state are governed by the Jacobian $D\mathbf{f}(\mathbf{x}_{LB})$, whose eigenvalues are $\lambda_1, \dots, \lambda_N \in \mathbb{C}$.

Let us assume that the eigenvalues satisfy the non-resonance conditions

$$\sum_{j=1}^N m_j \lambda_j \neq \lambda_i, \quad m_j \in \mathbb{N}, \quad (2.3)$$

for $i = 1, \dots, N$ and $\sum_{j=1}^N m_j \geq 2$. The conditions (2.3) guarantee the applicability of the linearisation theorem of Sternberg (1958) for class C^∞ dynamical systems, such as (2.2). This theorem states the existence of a smooth transformation mapping (2.2) to its linearisation around the edge state.

As shown by Haller *et al.* (2023), this implies the existence of a family of SSMs tangent to any given spectral subspace of the linearised dynamics. Precisely one of these SSMs, the primary SSM, is C^∞ smooth, whereas all others have reduced differentiability. We target the slowest family of SSMs, tangent to the slowest spectral subspace spanned by eigenvectors of the linearisation with eigenvalues closest to the imaginary axis. This is a normally attracting slow manifold if the remaining eigenvalues of $D\mathbf{f}(\mathbf{x}_{LB})$ have negative real parts.

The non-resonance conditions (2.3) are violated only if a resonance occurs among both the real and imaginary parts of the eigenvalues, i.e. the frequencies and the decay rates, simultaneously. As discussed by Haller *et al.* (2023), such an exact resonance among generic complex numbers is highly unlikely in a typical finite-dimensional system, such as (2.2). In addition, 1 : 1 resonances appearing as repeated eigenvalues enforced by physical symmetries do not violate (2.3). Therefore, we expect that a unique, C^∞ , mixed-mode SSM of the edge state exists. Alternatively, we can also invoke the results of Buza (2024), who showed that a C^1 -smooth pseudo-unstable manifold exists for the Navier–Stokes equations (2.1).




Figure 1. (a) Streamwise averaged velocity field of the edge state. The streamwise velocity is colour-coded, and the spanwise and wall-normal velocities are indicated as streamlines. (b) The spectrum of the edge state, computed by Channelflow. The eigenvalues associated with the spectral subspace E spanned by v_1 , $\text{Re } v_2$ and $\text{Im } v_2$, to which the SSM $\mathcal{W}(E)$ is tangent, are marked by red crosses. (c) Same as (a) for the vectors v_1 , $\text{Re } v_2$ and $\text{Im } v_2$. This figure is also available as a Jupyter notebook (<https://www.cambridge.org/S0022112025107957/JFM-Notebooks/files/figure1/figure1.ipynb>).

Liu *et al.* (2024) and Xu *et al.* (2024) showed that slow mixed-mode SSMs tangent to both stable and unstable linear modes can contain the chaotic attractor of a dynamical system. Therefore, they often serve as inertial manifolds (Foias *et al.* 1988).

Kreilos & Eckhardt (2012) constructed a one-dimensional (1-D) Poincaré map by tracking the maxima of the kinetic energy signal on the chaotic attractor. Although such a construction only decreases the dimension of the system (2.2) by one, this 1-D map already reveals periodic orbits in the attractor. These findings indicate that the attractor is likely low-dimensional.

3. Results

Figure 1(a,b) show the edge state with its spectrum. The unstable manifold is 1-D, and is tangent to the eigenvector v_1 . The next slowest eigenvalues are stable and form a complex conjugate pair, with the eigenvectors v_2 , $v_3 = \bar{v}_2$. The linear span of the slowest eigenvectors, corresponding to an unstable real eigenvalue and a pair of stable complex eigenvalues, is the slowest mixed-mode spectral subspace $E = \text{span}(v_1, \text{Re } v_2, \text{Im } v_2)$. These eigenvectors are shown in figure 1(c). The slow SSM, $\mathcal{W}(E)$, is therefore also 3-D and is constructed as a graph over the subspace spanned by these eigenvectors.

We follow the SSMLearn algorithm of Cenedese *et al.* (2022) to approximate the mixed-mode SSM, $\mathcal{W}(E)$, from data. Our training trajectories are initialised near the edge state, and are first attracted to the slow SSM before converging to either the laminar state or the chaotic attractor. The DNS solver Channelflow also computes the eigenvectors $v_{1,2,3}$ using the Arnoldi method, which we use to enforce the exact tangency between E and $\mathcal{W}(E)$, i.e. we parametrise the SSM over the subspace spanned by v_1 , $\text{Re } v_2$ and $\text{Im } v_2$.

The coordinate chart returning the reduced coordinates η is a projection to the eigenmodes

$$\eta = V^T x, \quad (3.1)$$

where V is the matrix containing v_1 , $\text{Re } v_2$ and $\text{Im } v_2$. In addition, the eigenvectors allow us to ensure that the training trajectories lie close to the SSM. Specifically, close to the fixed point, the SSM $\mathcal{W}(E)$ is well approximated by the spectral subspace E . We initialise the trajectories as small perturbations of the edge state along the spectral subspace, which




Figure 2. (a) The average kinetic energy along training trajectories. (b) Correlation dimension estimation based on (3.3), in the full phase space and the reduced phase space. The corresponding power-law fits are shown in blue and orange, respectively. (c) The reduced coordinates of the same trajectories as they converge to the chaotic attractor. The inset shows the laminarising trajectories as well. This figure is also available as a Jupyter notebook (<https://www.cambridge.org/S0022112025107957/JFM-Notebooks/files/figure2/figure2.ipynb>).

results in initial velocity fields of the form

$$\mathbf{x} = \mathbf{x}_{LB} + \alpha_1 \mathbf{v}_1 + \alpha_2 \operatorname{Re} \mathbf{v}_2 + \alpha_3 \operatorname{Im} \mathbf{v}_2, \quad (3.2)$$

with small coefficients $\alpha_{1,2,3}$. The sign of α_1 decides on which side of the edge of chaos the initial condition lies. Since we prepare initial conditions to lie close to the SSM, and we consider that the tangent space of the SSM is known *a priori*, our approach is not strictly data-driven, but data-assisted, using the terminology of Cenedese *et al.* (2025).

We advect 12 initial conditions (3.2) using Channelflow up to time $T_{max} = 2000$, sampling the trajectories at integer multiples of the dimensionless time unit $\Delta t = 1$. For comparison, we note that the characteristic time scale of the instability of the edge state is approximately $1/\lambda_1 = 21$, as seen in figure 1(c). In total, this results in $N_d = 24\,000$ data points in the training set. The kinetic energy $E = (1/2) \|\mathbf{u}\|^2$ of these trajectories, averaged over the domain Ω , is shown in figure 2(a). The reduced coordinates of the training trajectories, given by (3.1), are shown in figure 2(c). On the chaotic side of the edge, we see saddle-spiral type dynamics, reminiscent of a Rössler-type attractor (Rössler 1976), while on the other side, the dynamics are essentially 1-D, leading to rapid laminarisation.

To find further evidence of the low-dimensionality of the chaotic attractor, we estimate its correlation dimension using the training trajectories. As defined by Grassberger & Procaccia (1983), the correlation dimension is the scaling exponent γ in the relation

$$C(\varepsilon) \sim \varepsilon^\gamma \quad \text{as } \varepsilon \rightarrow 0, \quad (3.3)$$

where $C(\varepsilon)$ is the correlation sum, i.e. the number of pairs of points in the attractor that are separated by a distance less than $\varepsilon > 0$.

We computed the correlation sum $C(\varepsilon)$ in the full, N -dimensional phase space and in the reduced, 3-D phase space. The results are shown in figure 2(b). A linear fit on a logarithmic scale returns a correlation dimension of approximately $\gamma = 1.8$ in both cases.

Therefore, the chaotic attractor is indeed low-dimensional, and no topological information is lost by restricting the dynamics to the SSM.

We note that in principle, the chaotic attractor might not be smoothly embedded in three dimensions. This would be manifested by self-intersections of the trajectories in the reduced space. In this case, however, we could embed it in an SSM of the same edge state that is higher than 2γ -dimensional, as required by the Whitney (1944) embedding theorem. The reduced trajectories in figure 2(c) show no self-intersection, therefore we restrict the dynamics to the lowest dimensional slow SSM, $\mathcal{W}(E)$.

More generally, the SSM $\mathcal{W}(E)$ may develop a fold over its tangent space E , hence the parametrisation, constructed as a graph over E , may become singular. This would also be signalled by self-intersections of trajectories in the reduced space. Such examples are not common, but a notable one is discussed by Cenedese *et al.* (2022), who showed that the two-dimensional (2-D) SSM in the flow behind a cylinder folded over its tangent space. As in that case, reparametrising the manifold over a different set of observables generally solves this issue.

3.1. The geometry of the SSM

Having defined the coordinate chart (3.1), we proceed by defining the parametrisation of the SSM. We first shift the coordinates to place the edge state $\mathbf{x}_{LB}$ at the origin. We then seek the parametrisation in the polynomial form

$$\mathbf{x} = \mathbf{W}(\boldsymbol{\eta}) = \mathbf{V}\boldsymbol{\eta} + \sum_{n=2}^{M_p} \sum_{i+j+k=n} \mathbf{W}_{ijk} \eta_1^i \eta_2^j \eta_3^k, \quad (3.4)$$

where $\mathbf{W}_{ijk}$ are the coefficients of the monomial terms. In the expression for the parametrisation $\mathbf{W}(\boldsymbol{\eta})$, we explicitly distinguish the linear part as $\mathbf{D}\mathbf{W}(\mathbf{0}) = \mathbf{V}\boldsymbol{\eta}$. We use SSMLearn to determine the higher-order coefficients by minimising the error between the training data and (3.4). Specifically, we minimise the reconstruction error between $\mathbf{x}$ and $\mathbf{W}(\boldsymbol{\eta}) = \mathbf{W}(\mathbf{V}^T \mathbf{x})$ defined as

$$E_r = \sum_{j=1}^{N_d} \|\mathbf{x}^j - \mathbf{W}(\mathbf{V}^T \mathbf{x}^j)\|^2, \quad (3.5)$$

where $\mathbf{x}^j$ denotes the j th snapshot in the training set. Since we have prescribed the tangent space of the SSM, the minimisation problem is equivalent to linear regression and can be solved in closed form.

Evaluating the error (3.5) on a validation trajectory, we determine that the polynomial order $M_p = 5$ results in the optimal reconstruction error. For more information, we refer to the JFM Notebook accompanying figure 3.

3.2. The reduced dynamics on the SSM

After identifying the geometry of the SSM, $\mathcal{W}(E)$, we approximate the reduced dynamics. We model it as a discrete dynamical system given by the 3-D iterated map

$$\boldsymbol{\eta}_{n+1} = \mathbf{F}(\boldsymbol{\eta}_n), \quad \mathbf{F} : \mathbb{R}^3 \longrightarrow \mathbb{R}^3, \quad (3.6)$$

where $\boldsymbol{\eta}_n$ is the reduced state at time step n .

Although the flow of (2.2) is fundamentally continuous in time, we use a sampling time $\Delta t = 1$ to obtain the reduced map (3.6).




Figure 3. (a) Model predictions of test trajectories. (b) Power spectral densities computed from a chaotic kinetic energy signal in the full model and the SSM-reduced model. (c) A subset of the edge of chaos, constructed as the boundary of the basins of attraction in the reduced model, shown with the training trajectories. Trajectories initialised on either side of the edge of chaos are also indicated as black lines. This figure is also available as a Jupyter notebook (<https://www.cambridge.org/S0022112025107957/JFM-Notebooks/files/figure3/figure3.ipynb>).

A polynomial approximation of the reduced dynamics is insufficient for capturing the chaotic dynamics, as reported by Liu *et al.* (2024) and Xu *et al.* (2024). Instead, we use the SSMLearn algorithm with an alternative interpolation method. In particular, we approximate the dynamics as a linear combination of radial basis functions (RBFs) (Buhmann 2003). Interpolation using RBFs is a standard function approximation method, often used to interpolate unstructured data. In addition, computing the interpolants does not require iterative training.

The map (3.6) is approximated in the form

$$\mathbf{F}(\boldsymbol{\eta}) = \sum_{i=1}^{N_d} \mathbf{c}_i k(\|\boldsymbol{\eta}_i - \boldsymbol{\eta}\|), \quad (3.7)$$

where $\boldsymbol{\eta}_i$, $i = 1, \dots, N_d$, are the training data points, and $k: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is a radial function, where $\mathbb{R}_{\geq 0}$ denotes the non-negative reals. The coefficients $\mathbf{c}_i \in \mathbb{R}^3$ are determined by linear regression, minimising the squared error, analogously to (3.5), between the training data and the RBF approximation. We use the multiquadric kernel $k(r) = \sqrt{r^2 + \epsilon^2}$, which is a popular choice for RBF approximation (Fasshauer 2007); it is also a default option in the implementation that we follow, available from the `scipy` library (Virtanen *et al.* 2020). We select the scale parameter as $\epsilon = 10^{-8}$.

Since the RBF approximation interpolates the training trajectories, we need to use the reduced dynamics (3.7) to predict the time evolution of previously unseen initial conditions, i.e. test trajectories, to show that we avoid overfitting. We show two such predictions in figure 3(a). The initial conditions are close to the edge state, but on opposite sides of the edge of chaos. The laminarising trajectory is accurately predicted over the

whole time interval. As expected, predictions of the chaotic trajectory are accurate only on shorter time scales. This is due to the sensitivity of chaotic trajectories to initial conditions.

We can also compute relevant statistics on the chaotic attractor. In [figure 3\(b\)](#), we show the broad power spectral density obtained from the Fourier transform of the kinetic energy signal, characteristic of chaotic dynamics. A comparable time series of the full model yields a similar spectrum, with inaccuracies visible only towards high frequencies.

We further investigate the chaotic dynamics on the SSM by constructing the basins of the two coexisting attractors. We fill the domain of the reduced phase space shown in [figure 3\(c\)](#) with a total of 10^6 initial conditions, and iterate them forwards under the reduced dynamics (3.7) for $T = 1000$ steps. We then record the initial conditions of the laminarising trajectories. This yields a characteristic function of the basin of attraction of the laminar state, with the basin of the chaotic attractor obtained as its complement.

Instead of visualising the characteristic function directly in the reduced phase space, we construct the boundary between the two basins. We define this basin boundary as the level set of the characteristic function at value 0.5. The largest connected component of this level set, approximating the stable manifold of the edge state, is shown in [figure 3\(c\)](#). This set is a 2-D intersection of the edge of chaos with the SSM $\mathcal{W}(E)$.

We note, however, that since all our training trajectories have been initialised near the edge state, the basin boundary that we obtain is necessarily only a local approximation. Indeed, as shown by Kreilos & Eckhardt (2012), the basin of attraction of the chaotic attractor has a bubble-like shape, encircling the attractor. To accurately capture this in our reduced model, trajectories distributed along the entire basin boundary, also on both sides of the upper edge of chaos (Budanur *et al.* 2020), would need to be included in the training. This would increase the size of the required training set, and since our aim is not to map out the global shape of the basin boundary, we do not pursue this feature further.

The reduced dynamics (3.7) can also be used to compute the Lyapunov exponents (Oseledets 1968) of the chaotic attractor. The Lyapunov exponents are defined as the average exponential growth rates of perturbations along the chaotic attractor. Here, we compute the leading exponent, i.e. the most positive Lyapunov exponent, by tracking the average growth rate of perturbations in an ensemble of trajectories initialised close to an arbitrary point on the attractor (Cvitanović *et al.* 2016). Denoting the distance between two initially close trajectories as $\delta(t)$, the distance grows exponentially as $\delta(t) \sim e^{\Lambda t}$, where Λ is the leading Lyapunov exponent.

This computation can be carried out for the full system (Nastac *et al.* 2017), as well as for the SSM-reduced dynamics. [Figure 4\(a\)](#) shows the ensemble-averaged divergence of trajectories. By fitting a line to its initial, exponential trend, we find that the largest Lyapunov exponent is $\Lambda_{SSM} = 0.0150 \pm 0.0002$. This overestimates the leading Lyapunov exponent of the full system, which is $\Lambda_{DNS} = 0.0090 \pm 0.0001$. As a possible cause of this discrepancy, we note that our training set contains substantial data near the unstable edge state, whose eigenvalue $\lambda_1 = 0.047$ (see [figure 1\(b\)](#)) is considerably larger than Λ_{DNS} . In principle, this bias could be reduced by including more data from trajectories on the attractor.

We have also employed the recent neural ODE-based data-driven modelling method of Linot & Graham (2023) with our training data. The DManD algorithm reduces the dynamics to an inertial manifold by an initial linear projection to 500 proper orthogonal decomposition (POD) modes, followed by the application of an autoencoder. This maps the data to a latent space of dimension d_h , parametrising the inertial manifold. The reduced dynamics is then modelled as a neural ODE trained on the dynamics of the latent variables. We used the code published by Linot & Graham (2023) to train DManD models with




Figure 4. (a) Average rate of separation between nearby trajectories in the SSM-based model and in the DNS. (b) Relative reconstruction error of the autoencoders of the DManD models for various latent space dimensions d_h . (c,d) Average rates of separation of nearby trajectories in the best and worst DManD models, as defined in the text. This figure is also available as a Jupyter notebook (<https://www.cambridge.org/S0022112025107957/JFM-Notebooks/files/figure4/figure4.ipynb>).

$d_h = 3, \dots, 13$. Due to the stochastic nature of the training algorithm, we repeated the training ten times for each latent dimension, as described by Linot & Graham (2020).

Figure 4(b) shows the autoencoder reconstruction error, which saturates at approximately $d_h = 5$. Depending on the training process, there can be a considerable difference in the performance of models with otherwise identical hyperparameters. We compute the reconstruction error of all ten neural ODE models on a validation trajectory for each d_h . Based on this metric, we select the best and worst performing models.

To compare the DManD models to our SSM-based models, we first show $d_h = 3$, which matches the dimension of our SSM. As figure 4(c) shows, many of the models that we trained produce non-chaotic behaviour, and the best model underestimates the Lyapunov exponent. On the other hand, models with a higher-dimensional latent space achieve a very good representation of the chaotic dynamics on the attractor. We show the estimated leading Lyapunov exponent for $d_h = 13$ in figure 4(d).

We also observed that even if they perform well on chaotic trajectories, higher-dimensional DManD models tend to mispredict trajectories converging to the laminar attractor. This suggests that the DManD models have difficulties representing the two coexisting attractors of our system, which is to be expected since the method was originally proposed to model systems with a single global chaotic attractor. Addressing these issues will require more training data to resolve the basin boundary accurately, or adjustments to the network architecture. We have included further details of the training and evaluation of the DManD models in the JFM Notebook accompanying figure 4.

3.3. Poincaré map

Once the reduced dynamics (3.7) is obtained, we can reduce the model dimension even further by defining a Poincaré map (Guckenheimer & Holmes 1983). Since the sampling time $\Delta t = 1$ is small compared to other relevant time scales, we formally treat the trajectories of (3.6) as if they were continuous in time by linearly interpolating between successive discrete time steps. This allows us to compute the Poincaré map, which is defined for continuous-time dynamical systems.

The Poincaré (or first-return) map $\mathbf{P}(\boldsymbol{\eta})$ is then defined by successive intersections of the trajectory with a plane in the reduced phase space. We select the plane to be $\eta_2 = 0$, and record intersections with $(\eta_2)_{n+1} > 0 > (\eta_2)_n$. Figure 4(a) shows the intersection of the chaotic attractor with $\eta_2 = 0$.

Using the reduced dynamics, the map $\mathbf{P}$ can be evaluated arbitrarily many times to obtain a high-resolution representation of the chaotic attractor. This is shown in figure 4(b), which features a total of 10^6 intersection points generated by $\mathbf{P}$, revealing a much richer structure than the limited number of sample points obtained from the training trajectories. In particular, the SSM-reduced model revealed the small-scale fractal structure of the attractor in the inset of figure 5(b). Visually, the attractor resembles the attractor of the Hénon map, as remarked also by Kreilos & Eckhardt (2012).

The Poincaré section of the SSM, $\mathcal{W}(E)$, denoted $\mathcal{W}(E)|_{\eta_2=0}$, can be computed by simply setting $\eta_2 = 0$ in the parametrisation (3.4). In figure 5(c), we visualise the attractor of the Poincaré map on the SSM, by computing the kinetic energy of the flow fields.

We can also use the low-dimensional Poincaré map on the SSM to extract previously unobserved features of the chaotic dynamics. In particular, since the chaotic attractor is the closure of infinitely many unstable periodic orbits (Guckenheimer & Holmes 1983), we can use the Poincaré map to find some of these underlying orbits.

Computing periodic orbits by approximating the Poincaré map sampled using simulated trajectories has previously been applied to chaotic systems with low-dimensional attractors. Christiansen, Cvitanovic & Putkaradze (1997) computed periodic orbits of the Kuramoto–Shivashinsky equations (KSE) using a 1-D Poincaré map. A similar approach was used by Kreilos & Eckhardt (2012), who defined a Poincaré map by interpolation, using empirically computed successive maxima of the kinetic energy signal. More recently, Wang *et al.* (2025) used the cylinder torques measured in Taylor–Couette flow to define a 1-D map and prove its chaoticity. Also recently, Abadie *et al.* (2025) defined a Poincaré map in a low-dimensional subspace spanned by the leading POD modes and in the latent space of an autoencoder to find periodic orbits of the KSE.

With our SSM-reduced model, the Poincaré map is not approximated directly from full system simulations. Instead, it is a prediction of the reduced model, derived from the formally continuous time dynamics (3.6). As figure 5(b) shows, the training data used to obtain the continuous time reduced model (3.6) contain a much lower number of observed intersections with the Poincaré section. Although they do not define a Poincaré map, Linot & Graham (2023) also use the low-dimensional model to find periodic orbits of the full system.

The Poincaré map $\mathbf{P}$ bypasses the problem of finding an appropriate period, because periodic orbits appear as fixed points of iterates of $\mathbf{P}$. We therefore search for solutions of the equation $\mathbf{P}^n(\boldsymbol{\eta}) - \boldsymbol{\eta} = \mathbf{0}$, while systematically increasing n . We use a Newton–Raphson-like method to find these roots.

We initialise the root-finding method with $n = 1, \dots, 5$ from random initial conditions distributed along the attractor. Upon successful identification of a periodic orbit, the initial condition is then provided as an initial guess to Channelflow’s Newton–Krylov solver




Figure 5. (a) Black dots indicate the Poincaré section of the chaotic attractor with the $\eta_2 = 0$ plane (grey). (b) Black dots indicate intersections computed based on the training trajectories. Red dots are iterations of the SSM-reduced Poincaré map. (c) Poincaré section of the SSM $\mathcal{W}(E)$, containing the chaotic attractor (black) and periodic orbits (red). A period-3 orbit is highlighted in white, with the flow fields shown in (d–f) using the same visualisation as in figure 1. The kinetic energy of this orbit is shown in (g). This figure is also available as a Jupyter notebook (<https://www.cambridge.org/S0022112025107957/JFM-Notebooks/files/figure5/figure5.ipynb>).

(Viswanath 2007). Although we did not include any periodic orbits in the training data, we successfully found eight periodic ECSs of the full system using the initial guesses from the reduced model. One of these orbits is shown in figure 5(d–g).

We emphasise that our goal was not to compile an extensive library of periodic orbits of this particular system, but to demonstrate that an SSM-based reduced model already captures the essential chaotic dynamics. Indeed, the reduced dynamics enables the efficient sampling of the natural measure of the chaotic attractor to compute the probability distribution of quantities of interest, such as the spectrum of the kinetic energy in figure 3(b).

4. Conclusions and discussion

We have applied the theory of spectral submanifolds (SSMs) to plane Couette flow in the parameter regime supporting a chaotic attractor. We restricted the flow to the slowest 3-D mixed-mode SSM of the edge state, which functions as an inertial manifold. Our very-low-dimensional reduced model is given by an iterated map interpolated using radial basis functions. Although contemporary machine learning methods, such as neural ODEs

(Linot & Graham 2020), could also be implemented to model the dynamics, we found that simpler interpolation methods already deliver excellent accuracy.

Compared to these machine learning methods, the classical function approximation used here has the advantage that it does not require any iterative training. This makes the models less data-hungry and faster to compute, and their performance is more predictable. This is due to the exact mathematical background supporting the existence and smoothness properties of the underlying SSMs. In addition, we found that the state-of-the-art neural-network-based model DManD (Linot & Graham 2023) may have difficulties modelling the dynamics due to the coexisting chaotic and laminar attractors. In order for the neural ODE model to produce chaotic dynamics consistently, larger-dimensional inertial manifolds were required.

We have also defined a 2-D Poincaré map associated with the SSM-reduced model to extract previously unseen characteristics of the dynamics. In particular, the Poincaré map revealed the fractal structure of the chaotic attractor, and facilitated the search for periodic orbits. Although Kreilos & Eckhardt (2012) inferred a similar 1-D map based on simulated time histories of the kinetic energy to find periodic orbits, our approach makes a rigorous connection between the dynamics in the full phase space and a 2-D Poincaré section in the reduced phase space.

Our SSM-reduced model has been proven to generalise effectively and represent the chaotic attractor accurately. Therefore, we believe that constructing SSM models, attached to known ECSs, is a promising method for modelling other flows, too. In particular, they allow efficient sampling of the probability distributions of physical variables without relying on periodic orbit theory (Cvitanović *et al.* 2016).

Here, we have focused on plane Couette flow in a minimal flow unit at a relatively low Reynolds number. This allowed us to conclude the existence of an SSM serving as an inertial manifold, since a hyperbolic ECS close to the chaotic attractor, with an slow attracting mixed-mode SSM, was available. As a result, although the system exhibits temporal chaos, the observed spatial flow structures remain simple. Nevertheless, successful modelling of such simple, but physically relevant systems is an important step towards modelling flows with a higher degree of spatial complexity.

A logical future direction will be to pursue a similar model reduction approach in the same system at a higher Reynolds number, after the boundary crisis (Kreilos *et al.* 2014). For example, another popular parameter choice is $Re = 400$, where the lower branch fixed point persists with a 1-D unstable manifold. The flow is more complex at this Reynolds number, and the dimension of its chaotic set likely increases but stays below 18, according to Linot & Graham (2023). The main difficulty in that case comes from the transient nature of the chaotic dynamics. Since this dynamics is supported on a chaotic saddle, trajectories may need to be carefully initialised near its stable manifold in order to explore its dynamics for sufficiently long times.

We finally note that systems exhibiting a supercritical transition to chaos and turbulence may also be reduced to a low-dimensional SSM. In that case, the model needs to be anchored to the unstable laminar state, which is often available in closed form. The appropriate slowest mixed-mode SSM of sufficiently high dimension can then be constructed in a way similar to that presented here. We expect an SSM-reduced model to be successful in those cases as well, at least for moderately chaotic flows. This is also supported by the recent results of Wang *et al.* (2025), who constructed an empirical 1-D Poincaré map for a low Reynolds number Taylor–Couette flow and found unstable periodic orbits embedded in the attractor.

Supplementary material. Computational Notebook files are available as supplementary material at <https://doi.org/10.1017/jfm.2025.107954> and online at <https://www.cambridge.org/S0022112025107957/JFM-Notebooks>.

Acknowledgements. We are grateful to A. Linot and M. Graham for making their DManD code available, and for helpful discussions and useful advice on its implementation. The numerical simulations were performed on the Euler cluster operated by the High Performance Computing group at ETH Zürich.

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

REFERENCES

- ABADIE, M., BECK, P., PARKER, J.P. & SCHNEIDER, T.M. 2025 The topology of a chaotic attractor in the Kuramoto–Sivashinsky equation. *Chaos: Interdisciplinary J. Nonlinear Sci.* **35** (1), 013123.
- AHMED, O., TENNIE, F. & MAGRI, L. 2025 Robust quantum reservoir computers for forecasting chaotic dynamics: generalized synchronization and stability, arXiv preprint arXiv: 2506.22335.
- AVILA, M., BARKLEY, D. & HOF, B. 2023 Transition to turbulence in pipe flow. *Annu. Rev. Fluid Mech.* **55** (1), 575–602.
- BUDANUR, N.B., MARENSI, E., WILLIS, A.P. & HOF, B. 2020 Upper edge of chaos and the energetics of transition in pipe flow. *Phys. Rev. Fluids* **5**, 023903.
- 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.
- BUHMANN, M.D. 2003 *Radial Basis Functions: Theory and Implementations*. Cambridge University Press.
- BUZA, G. 2024 Spectral submanifolds of the Navier–Stokes equations. *SIAM J. Appl. Dyn. Syst.* **22** (2), 1052–1089.
- CAVALIERI, A.V.G. & NOGUEIRA, P.A.S. 2022 Reduced-order Galerkin models of plane Couette flow. *Phys. Rev. Fluids* **7**, L102601.
- CENEDESE, M., AXÅS, J., BÄUERLEIN, B., AVILA, K. & HALLER, G. 2022 Data-driven modeling and prediction of non-linearizable dynamics via spectral submanifolds. *Nat. Commun.* **13** (1), 872.
- CENEDESE, M., MARCONI, J., HALLER, G. & JAIN, S. 2025 Data-assisted non-intrusive model reduction for forced nonlinear finite elements models. *Nonlinear Dyn.* **113** (7), 6465–6489.
- CHRISTIANSEN, F., CVITANOVIĆ, P. & PUTKARADZE, V. 1997 Spatiotemporal chaos in terms of unstable recurrent patterns. *Nonlinearity* **10** (1), 55.
- CVITANOVIĆ, P., ARTUSO, R., MAINIERI, R., TANNER, G. & VATTAY, G. 2016 *Chaos: Classical and Quantum*. Niels Bohr Institute.
- ECKHARDT, B. 2018 Transition to turbulence in shear flows. *Physica A: Stat. Mech. Appl.* **504**, 121–129.
- ECKMANN, J.P. 1981 Roads to turbulence in dissipative dynamical systems. *Rev. Mod. Phys.* **53**, 643–654.
- FAISST, H. & ECKHARDT, B. 2003 Traveling waves in pipe flow. *Phys. Rev. Lett.* **91** (22), 224502.
- FASSHAUER, G.E. 2007 *Meshfree Approximation Methods with Matlab*. World Scientific.
- FOIAS, C., SELL, G.R. & TEMAM, R. 1988 Inertial manifolds for nonlinear evolutionary equations. *J. Differ. Equ.* **73** (2), 309–353.
- GIBSON, J.F. 2014 Channelflow: a spectral Navier–Stokes simulator in C++. *Tech. Rep.* University of New Hampshire, Channelflow.org.
- GIBSON, J.F., *et al.* 2019 Channelflow 2.0. *Tech. Rep.* manuscript in preparation, channelflow.ch
- GIBSON, J.F., HALCROW, J. & CVITANOVIĆ, P. 2008 Visualizing the geometry of state space in plane Couette flow. *J. Fluid Mech.* **611**, 107–130.
- GIBSON, J.F., HALCROW, J. & CVITANOVIĆ, P. 2009 Equilibrium and travelling-wave solutions of plane Couette flow. *J. Fluid Mech.* **638**, 243–266.
- GOLLUB, J.P. & SWINNEY, H.L. 1975 Onset of turbulence in a rotating fluid. *Phys. Rev. Lett.* **35**, 927–930.
- 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.
- GRASSBERGER, P. & PROCACCIA, I. 1983 Characterization of strange attractors. *Phys. Rev. Lett.* **50**, 346–349.
- GREBOGI, C., OTT, E. & YORKE, J.A. 1983 Crises, sudden changes in chaotic attractors, and transient chaos. *Physica D: Nonlinear Phenom.* **7** (1), 181–200.
- GUCKENHEIMER, J. & HOLMES, P. 1983 Nonlinear oscillations, dynamical systems, and bifurcations of vector fields. In *Applied Mathematical Sciences*, vol. 42.
- HALLER, G. 2025 *Modeling Nonlinear Dynamics from Equations and Data*. Springer.
- HALLER, G., KASZÁS, B., LIU, A. & AXÅS, J. 2023 Nonlinear model reduction to fractional and mixed-mode spectral submanifolds. *Chaos* **33**, 063138.

- JAEGER, H. 2001 The ‘echo state’ approach to analysing and training recurrent neural networks – with an erratum note. GMD Report 148. German National Research Center for Information Technology (GMD).
- KASZÁS, B., CENEDESE, M. & HALLER, G. 2022 Dynamics-based machine learning of transitions in Couette flow. *Phys. Rev. Fluids* **7** (8), L082402.
- KASZÁS, B. & HALLER, G. 2024 Capturing the edge of chaos as a spectral submanifold in pipe flows. *J. Fluid Mech.* **979**, A48.
- KAWAHARA, G. & KIDA, S. 2001 Periodic motion embedded in plane Couette turbulence: regeneration cycle and burst. *J. Fluid Mech.* **449**, 291–300.
- KLEISER, L. & SCHUMANN, U. 1980 *Treatment of Incompressibility and Boundary Conditions in 3-D Numerical Spectral Simulations of Plane Channel Flows*. Vieweg+Teubner Verlag.
- KREILOS, T. & ECKHARDT, B. 2012 Periodic orbits near onset of chaos in plane Couette flow. *Chaos: Interdisciplinary J. Nonlinear Sci.* **22** (4), 047505.
- KREILOS, T., ECKHARDT, B. & SCHNEIDER, T.M. 2014 Increasing lifetimes and the growing saddles of shear flow turbulence. *Phys. Rev. Lett.* **112**, 044503.
- LAI, Y.-C. & TÉL, T. 2011 Transient chaos. In *Applied Mathematical Sciences*, vol. 173.
- LINOT, A.J. & GRAHAM, M.D. 2020 Deep learning to discover and predict dynamics on an inertial manifold. *Phys. Rev. E* **101** (6), 062209.
- LINOT, A.J. & GRAHAM, M.D. 2023 Dynamics of a data-driven low-dimensional model of turbulent minimal Couette flow. *J. Fluid Mech.* **973**, A42.
- LIU, A., AXÅS, J. & HALLER, G. 2024 Data-driven modeling and forecasting of chaotic dynamics on inertial manifolds constructed as spectral submanifolds. *Chaos* **34** (3), 033140.
- MARGAZOGLU, G. & MAGRI, L. 2023 Stability analysis of chaotic systems from data. *Nonlinear Dyn.* **111** (9), 8799–8819.
- MOORE, D.R., TOOMRE, J., KNOBLOCH, E. & WEISS, N.O. 1983 Period doubling and chaos in partial differential equations for thermosolutal convection. *Nature* **303** (5919), 663–667.
- NAGATA, M. 1990 Three-dimensional finite-amplitude solutions in plane Couette flow: bifurcation from infinity. *J. Fluid Mech.* **217**, 519–527.
- NASTAC, G., LABAHN, J.W., MAGRI, L. & IHME, M. 2017 Lyapunov exponent as a metric for assessing the dynamic content and predictability of large-eddy simulations. *Phys. Rev. Fluids* **2**, 094606.
- OSELEDETS, V.I. 1968 A multiplicative ergodic theorem. Lyapunov characteristic numbers for dynamical systems. *Trans. Moscow Math. Soc.* **19**, 197–231.
- ÖZALP, E. & MAGRI, L. 2025 Stability analysis of chaotic systems in latent spaces. *Nonlinear Dyn.* **113**, 13791–13806.
- PAGE, J. & KERSWELL, R.R. 2019 Koopman mode expansions between simple invariant solutions. *J. Fluid Mech.* **879**, 1–27.
- PAGE, J., NORGAAARD, P., BRENNER, M.P. & KERSWELL, R.R. 2024 Recurrent flow patterns as a basis for two-dimensional turbulence: predicting statistics from structures. *Proc. Natl Acad. Sci.* **121** (23), e2320007121.
- PATHAK, J., HUNT, B., GIRVAN, M., LU, Z. & OTT, E. 2018 Model-free prediction of large spatiotemporally chaotic systems from data: a reservoir computing approach. *Phys. Rev. Lett.* **120**, 024102.
- PATHAK, J., LU, Z., HUNT, B.R., GIRVAN, M. & OTT, E. 2017 Using machine learning to replicate chaotic attractors and calculate Lyapunov exponents from data. *Chaos: Interdisciplinary J. Nonlinear Sci.* **27** (12), 121102.
- POMEAU, Y. & MANNEVILLE, P. 1980 Intermittent transition to turbulence in dissipative dynamical systems. *Commun. Math. Phys.* **74**, 189–197.
- RACCA, A., DOAN, N.A.K. & MAGRI, L. 2023 Predicting turbulent dynamics with the convolutional autoencoder echo state network. *J. Fluid Mech.* **975**, A2.
- RÖSSLER, O. 1976 An equation for continuous chaos. *Phys. Lett. A* **57** (5), 397–398.
- RUELLE, D. & TAKENS, F. 1971 On the nature of turbulence. *Commun. Math. Phys.* **20** (3), 167–192.
- SCHMID, P.J. 2010 Dynamic mode decomposition of numerical and experimental data. *J. Fluid Mech.* **656**, 5–28.
- SCHNEIDER, T.M., GIBSON, J.F., LAGHA, M., DE LILLO, F. & ECKHARDT, B. 2008 Laminar–turbulent boundary in plane Couette flow. *Phys. Rev. E* **78** (3), 037301.
- SKUFCA, J.D., YORKE, J.A. & ECKHARDT, B. 2006 Edge of chaos in a parallel shear flow. *Phys. Rev. Lett.* **96** (17), 174101.
- STERNBERG, S. 1958 On the structure of local homeomorphisms of Euclidean n -space, II. *Am. J. Maths* **80** (3), 623.
- SURI, B., PALLANTLA, R.K., SCHATZ, M.F. & GRIGORIEV, R.O. 2019 Heteroclinic and homoclinic connections in a Kolmogorov-like flow. *Phys. Rev. E* **100** (1), 013112.

- VIRTANEN, P. *et al.* 2020 SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Meth.* **17**, 261–272.
- VISWANATH, D. 2007 Recurrent motions within plane Couette turbulence. *J. Fluid Mech.* **580**, 339–358.
- WALEFFE, F. 2001 Exact coherent structures in channel flow. *J. Fluid Mech.* **435**, 93–102.
- WANG, B., AYATS, R., DEGUCHI, K., MESEGUER, A. & MELLIBOVSKY, F. 2025 Mathematically established chaos and forecast of statistics with recurrent patterns in Taylor–Couette flow. *J. Fluid Mech.* **1011**, R2.
- WANG, J., GIBSON, J. & WALEFFE, F. 2007 Lower branch coherent states in shear flows: transition and control. *Phys. Rev. Lett.* **98** (20), 204501.
- WEDIN, H. & KERSWELL, R.R. 2004 Exact coherent structures in pipe flow: travelling wave solutions. *J. Fluid Mech.* **508**, 333–371.
- WHITNEY, H. 1944 The self-intersections of a smooth n -manifold in $2n$ -space. *Ann. Maths* **45** (2), 220–246.
- WILLIS, A.P., CVITANOVIĆ, P. & AVILA, M. 2013 Revealing the state space of turbulent pipe flow by symmetry reduction. *J. Fluid Mech.* **721**, 514–540.
- XU, Z., KASZÁS, B., CENEDESE, M., BERTI, G., COLETTI, F. & HALLER, G. 2024 Data-driven modelling of the regular and chaotic dynamics of an inverted flag from experiments. *J. Fluid Mech.* **987**, R7.