Research Square

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

A wall-based, interpretable reduced-order model for real-time estimation of near-wall turbulence

Miguel Pérez Cuadrado

miguel.perez-cuadrado@citystgeorges.ac.uk

City St George's, University of London

Giorgio Maria Cavallazzi

City St George's, University of London

Andrea Nóvoa

Imperial College London

Alfredo Pinelli

City St George's, University of London

Research Article

Keywords: Wall turbulence, non-linear reduced-order modelling, radial basis functions, clustering, data assimilation, Kalman filter

Posted Date: August 4th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10550785/v1

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

Additional Declarations: No competing interests reported.

A wall-based, interpretable reduced-order model for real-time estimation of near-wall turbulence

M. Pérez Cuadrado[1*], G.M. Cavallazzi[1], A. Nóvoa[2,3], A. Pinelli[1]

[1*]Department of Engineering, School of Science and Technology, City St George's, University of London, London, UK.
[2]Department of Aeronautics, Imperial College London, South Kensington Campus, London, SW7 2BX, UK.
[3] I-X Centre for AI in Science, Imperial College London, London, SW7 2BX, UK.

*Corresponding author(s). E-mail(s): miguel.perez-cuadrado@citystgeorges.ac.uk;

Abstract

The near-wall region of a turbulent flow maintains a self-sustaining cycle of streaks and quasi-streamwise vortices, and this cycle leaves an uneven signature on the wall shear stress. We ask how much of the buffer layer dynamics is encoded in the time-evolving wall shear stress, or whether the wall stress is merely a passive footprint. To answer this question, we model the underlying dynamical system directly from wall measurements. Using snapshot data from a minimal turbulent channel at $\boldsymbol{Re_\tau = 180}$, we extract the proper orthogonal decomposition (POD) modes of the streamwise and spanwise wall shear stresses and approximate their attractor using radial basis functions (RBFs) with trajectory-adapted kernel widths. This yields a differentiable latent vector field that we can interrogate directly. Without any imposed supervision, the proposed ROM naturally separates the quiescent, vortex-dominated, and bursting phases of the wall cycle. In unforced long-term rollouts, the ROM trajectories stay bounded and recover the invariant measure. Since the RBF representation is differentiable, we are able to locate weakly unstable near-equilibrium points and resolve the non-normal growth and phase-volume contraction of the attractor. Finally, we assimilate sparse, noisy wall measurements with an ensemble Kalman filter, and the ROM then tracks the true flow trajectory. This demonstrates how tightly instantaneous wall shear constrains the cycle: the model recovers the invariant measure, and forecast errors driven by unobserved interior motions remain bounded and localised within the active phase.

Keywords: Wall turbulence, non-linear reduced-order modelling, radial basis functions, clustering, data assimilation, Kalman filter

1 Introduction

The dynamics of near-wall turbulence are organised around a self-sustaining cycle (Waleffe 1997). Within this cycle, elongated low-speed streaks form by lift-up (Brandt 2014), undergo a sinuous instability, break down (Andersson et al. 2001), and regenerate over a characteristic time of about a hundred wall units (Hamilton et al. 1995). The cycle is spatially compact, confined to a minimal flow unit (Jiménez and Moin 1991), and it remains autonomous when the outer flow is removed (Jiménez and Pinelli 1999).

The discovery of this regenerating cycle inspired active control schemes aimed at drag reduction. These approaches rely on sensing the velocity close to the wall and acting at the surface to manipulate the cycle, ranging from opposition control (Choi et al. 1994) to machine-learned policies (Guastoni et al. 2023; Cavallazzi et al. 2025). However, off-wall measurements required are impractical in experiments and at high Reynolds number, whereas the shear stress at the wall itself is more easily measurable (Naughton and Sheplak 2002; Löfdahl and Gad-el Hak 1999). This raises a fundamental question: can the dynamics of the near-wall cycle be fully reconstructed from wall-based variables alone? Inferring the dynamics of near-wall coherent structures from surface measurements is a long-standing objective, particularly in

wall-modelled large-eddy simulation (Piomelli and Balaras 2002; Bose and Park 2018). Since coherent motions in the buffer layer (Robinson 1991) and larger log-layer structures (Encinar and Jiménez 2019) imprint distinct signatures on wall pressure (Naguib et al. 2001) and wall shear stress (Kravchenko et al. 1993), the near-wall velocity field is, in principle, observable from wall-bounded quantities alone.

The wall footprint is variable and non-uniform across the cycle evolution. The near-wall velocity streaks leave a strong, direct signature in the streamwise wall shear stress (Robinson 1991), whereas the quasi-streamwise vortices are most cleanly identified through their imprint on the spanwise shear (Kravchenko et al. 1993). Despite this variability, the footprint still uniquely identifies the phase within the cycle. This enables static reconstructions of the instantaneous velocity field at a given wall-normal height (Güemes et al. 2021; Cuéllar et al. 2024), as well as dynamic estimations capable of providing time-dependent boundary conditions for online flow control in DNS (Cavallazzi et al. 2026).

Such a reconstruction is a snapshot, though, not the cycle: it fixes the state at one instant, and says nothing about whether the dynamics can be recovered from the wall. Recovering these dynamics, and therefore testing whether the wall retrieves the cycle at all, needs more than a snapshot: a model that advances the wall state on its own, staying bounded and reproducing the correct statistics under free integration.

However, even when surface measurements are used dynamically for real-time control, such estimations are state estimators for an external flow solver rather than representing an autonomous dynamical system of the cycle itself. Testing whether surface quantities fully constrain the underlying attractor requires going beyond estimator-driven integration to construct a standalone model. Such a model must advance the latent wall state on its own, remain bounded over long rollouts, and recover the true flow statistics under free integration. Capturing these nonlinear, long-term dynamics from data naturally motivates a reduced-order modelling (ROM) framework. By projecting the flow onto a compact latent manifold, which reduces the state dimension from millions of grid points in direct numerical simulations to a few latent variables, a ROM efficiently evolves the temporal state forward in time. Such models generally fall into intrusive and non-intrusive paradigms (Benner et al. 2015).

Intrusive models rely on explicit knowledge of the governing equations, projecting them onto the latent space. POD–Galerkin models are the canonical example: proper orthogonal decomposition (POD) (Lumley 1967; Sirovich 1987) supplies the compression, and the Navier–Stokes equations are projected onto the retained modes, globally (Cavalieri and Nogueira 2022) or on locally defined bases (Colanera and Magri 2025, 2026). Since the truncated small scales drain energy from the retained modes, such projections are prone to instability (Rempfer 2000), and the discarded scales must be reintroduced through a calibrated eddy-viscosity dissipation to recover stable long-time statistics (Aubry et al. 1988).

Non-intrusive models instead learn the reduced dynamics directly from data, absorbing the truncation effect implicitly and requiring no access to the governing operator. Some remain opaque, such as recurrent networks that predict accurately but act as black boxes (Srinivasan et al. 2019); others are interpretable, such as sparse symbolic regression (Brunton et al. 2016) or cluster-based networks of coherent states (Kaiser et al. 2014; Fernex et al. 2021).

An intrusive projection approach cannot be applied in this context. While the wall-shear dynamics are governed by the underlying Navier–Stokes equations, formulating a closed evolution equation strictly in terms of surface quantities requires unobserved off-wall flow fields. The reduced-order model must therefore be learned non-intrusively from data. Evaluating whether the near-wall cycle can be recovered strictly from surface observations imposes two simultaneous requirements. The model must govern autonomous temporal dynamics, and it must operate exclusively on wall-based inputs. Existing data-driven approaches address these requirements only in isolation. They either reconstruct instantaneous fields from wall measurements without capturing temporal evolution (Guastoni et al. 2021), or they model full-field temporal dynamics using off-wall information (Nakamura et al. 2021). A reduced-order framework that unifies both by evolving the dynamics of the cycle solely from surface quantities has yet to be established.

The model proposed here addresses this gap. Evaluating its autonomous rollouts demonstrates whether the wall-based dynamics successfully reproduce the statistical properties of the cycle. To further test whether surface measurements contain sufficient information to track the true state trajectory, the model forecast must be continuously confronted with online wall observations. Sequential data assimilation provides the necessary framework by combining state reconstruction with dynamic forecasting in real time. Specifically, an ensemble Kalman filter (EnKF)(Nóvoa et al. 2024; Nóvoa and Magri 2025) can correct the forecast using surface measurements, extending state estimation approaches previously demonstrated using the full Navier–Stokes solver (Colburn et al. 2011) or full-field data-driven models (Özalp et al. 2026).

The same filter has two uses: it pulls the model towards the true trajectory, and the mismatch between its forecast and the observations measures how much correction the model needs to stay consistent with

the data (Desroziers et al. 2005; Anderson 2007). We deploy this framework to evaluate whether wall-shear measurements alone can sustain and constrain the near-wall cycle. The low-dimensional state space is constructed using the non-intrusive manifold approach of Pérez Cuadrado et al. (2026) for a minimal turbulent channel (Jiménez and Moin 1991) at $Re_\tau = 180$ (Kim et al. 1987). By combining autonomous latent-space evolution with ensemble Kalman filtering, we assess both the long-term statistical fidelity and the short-term tracking capability of the wall-driven dynamics.

Unlike prior wall-based estimators that rely either on full governing equations (Colburn et al. 2011) or opaque data-driven architectures (Özalp et al. 2026), the present formulation is non-intrusive and interpretable. Since it recovers the invariant measure of the wall shear stress, the model is a well-posed forecast operator within the ensemble filter framework (Evensen 2003).

Using this approach, we show three capabilities of the proposed wall-driven framework. First, clustering the wall-shear trajectory maps the phases of the coherent-structure cycle, capturing the transition on the streak-vortex plane from a quiescent state to vortex dominance and mature streaks, while autonomous integration of the fitted vector field recovers the invariant measure of the skin friction. Second, dynamical analysis of the reduced field identifies regions of low phase velocity within the quiescent phase and isolates where disturbance growth peaks near streak breakdown. Third, ensemble Kalman filtering quantifies the degree of closure provided by surface measurements alone, using the required additive inflation to measure the unobserved interior dynamics. This framework therefore distinguishes the statistical autonomy of the wall shear stress from its deterministic tracking capability during real-time data assimilation.

The paper is structured around these methodological steps. Section 2 describes the turbulent channel simulation setup and the proper orthogonal decomposition of the wall shear stresses. Section 3 clusters the latent trajectory and demonstrates that the resulting partition recovers the phase-space topology of the near-wall cycle. Section 4 details the construction of the non-intrusive vector field and evaluates its invariant measure and forecast horizon during autonomous rollouts. Section 5 analyses the explicit vector field as a dynamical system, identifying regions of low phase velocity and the progression of non-normal growth around streak breakdown. Finally, Section 6 incorporates sparse, noisy surface observations via an ensemble Kalman filter, using the required additive inflation to quantify the unobserved interior dynamics.

2 Flow configuration and the wall observable

The wall shear stress data used to construct and evaluate the reduced model are generated by direct numerical simulation (DNS) of a turbulent channel flow at friction Reynolds number $Re_\tau = 180$, computed using the solver CaNS (Costa 2018). The solver integrates the incompressible Navier–Stokes equations on a staggered grid using second-order spatial finite differences. Time advancement is performed via a fractional-step method that combines a low-storage, three-stage Runge–Kutta scheme with a pressure-projection step at each substep to enforce solenoidal velocity fields. The flow is driven at a constant mass flow rate. We adopt the minimal channel configuration of Jiménez and Moin (1991), defining a half-channel domain with streamwise, spanwise, and wall-normal extents of $L_x^+ \times L_y^+ \times L_z^+ = 480 \times 140 \times 180$, capped by a midplane symmetry condition. The wall-parallel directions (x streamwise and y spanwise) are discretised using a uniform 64×64 grid, yielding spatial resolutions of $\Delta x^+ \approx 7.5$ and $\Delta y^+ \approx 2.2$. In the wall-normal direction, the grid is stretched to achieve a minimum near-wall resolution of $\Delta z^+_{\min} = 0.54$, consistent with established channel simulations (Kim et al. 1987). Time advancement proceeds with a constant step size of $\Delta t = 0.01$, and the streamwise (τ_x) and spanwise (τ_y) wall shear stress components are recorded every five time steps ($\Delta t^+ \approx 0.5$, i.e. half a viscous time unit apart). This temporal resolution is dictated by the accuracy requirements of the reduced description rather than storage constraints: the snapshot spacing must be small enough to yield smooth, accurate time derivatives for the target vector field while resolving the temporal evolution of the projected coefficients. Sampling continues until 10^5 snapshots are acquired, spanning roughly 5×10^4 viscous time units.

At each sampling instant, the two components (τ_x, τ_y) are stacked into a single vector $\boldsymbol{u} \in \mathbb{R}^N$ of dimension $N = 2 \times 64 \times 64 = 8192$. The $M = 10^5$ recorded fields are reduced by proper orthogonal decomposition (Lumley 1967; Sirovich 1987). The temporal mean is removed first, so that the basis describes the fluctuations about the mean wall state,

$$\bar{\boldsymbol{u}} = \frac{1}{M} \sum_{m=1}^{M} \boldsymbol{u}_m, \qquad \boldsymbol{u}'_m = \boldsymbol{u}_m - \bar{\boldsymbol{u}}. \tag{1}$$

Collecting the mean-subtracted fields as columns of the snapshot matrix $\mathbf{X} \in \mathbb{R}^{N\times M}$, its singular value decomposition (SVD) yields

$$\mathbf{X} = [\,\mathbf{u}'_1, \ldots, \mathbf{u}'_M\,] = \mathbf{\Phi}\mathbf{\Sigma}\mathbf{\Psi}^\top, \tag{2}$$

where the orthonormal columns of $\mathbf{\Phi}$ define the spatial POD modes, ordered by the singular values along the diagonal of $\mathbf{\Sigma}$. The squared singular values measure the variance of the wall shear-stress fluctuations captured by each mode. Retaining the leading d_o modes in $\mathbf{\Phi}_{d_o}$, the reduced state vector $\mathbf{a}(t)$ and the reconstructed wall-shear field are defined as

$$\mathbf{a}(t) = \mathbf{\Phi}_{d_o}^\top\big(\mathbf{u}(t) - \bar{\mathbf{u}}\big) \in \mathbb{R}^{d_o}, \qquad \mathbf{u}(t) \approx \bar{\mathbf{u}} + \mathbf{\Phi}_{d_o}\mathbf{a}(t). \tag{3}$$

Projecting each snapshot then generates the reduced state trajectory $\mathbf{a}_m = \mathbf{\Phi}_{d_o}^\top \mathbf{u}'_m$, upon which the model operates.

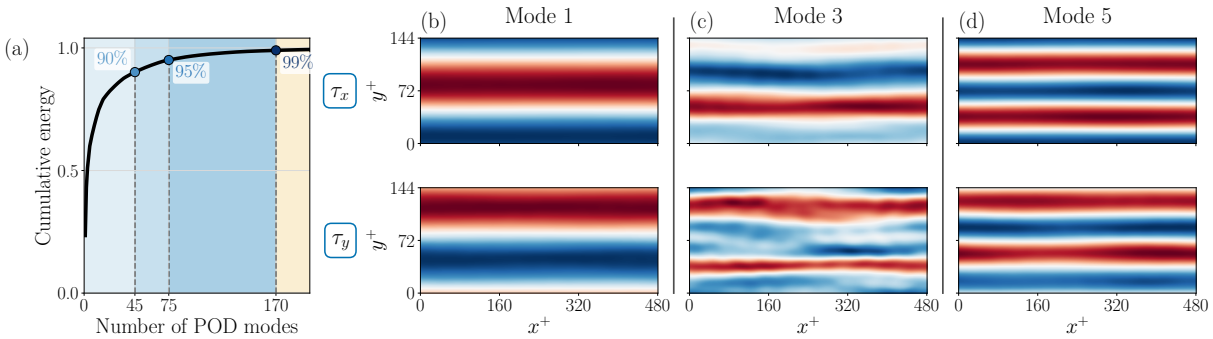


Fig. 1 Proper orthogonal decomposition of the wall shear stress. (a) Cumulative fluctuation energy against the number of retained modes, with the 90%, 95% and 99% thresholds reached at $d_o = 45$, 75 and 170 modes. (b–d) Streamwise τ_x (top row) and spanwise τ_y (bottom row) components of proper orthogonal modes 1, 3 and 5

We retain $d_o = 45$ modes, which capture approximately 90% of the wall shear-stress variance (Fig. 1a). The discarded tail consists predominantly of small-scale, low-amplitude wall-gradient fluctuations. Therefore, the choice of operating dimension targets the dynamically relevant scales rather than reconstruction accuracy alone. The fact that $d_o = 45$ adequately captures the dynamics of the near-wall cycle is supported by the autonomous integration statistics in Section 4, where expanding the basis up to $d_o = 75$ leaves the recovered climate unchanged. The clustering diagnostics in Section 3 independently exhibit the same elbow behaviour at $d_o = 45$, 75, and 170. For comparison, these higher-dimensional truncations are evaluated alongside the baseline 45-mode model.

The leading POD modes isolate the surface signature of the near-wall streaks, as shown by the streamwise and spanwise components of mode 1 in Fig. 1(b). Modes 1 and 2 form a quadrature pair, shifted in spatial phase by $\pi/2$, which together allow the streak structure to translate continuously along the spanwise direction. This behaviour is consistent with velocity-field decompositions of channel turbulence (Moin and Moser 1989; Aubry et al. 1988), whose spatial symmetries are preserved in the wall shear-stress field.

The evolution of the reduced coordinates (see Fig. 2) is modelled as an autonomous system of ordinary differential equations, $\dot{\mathbf{a}} = \mathbf{f}(\mathbf{a})$, where the non-linear right-hand side $\mathbf{f}$ is discretised using radial basis functions (RBFs). Fitting this vector field requires pairing each reduced state with its temporal derivative, $(\mathbf{a}_m, \dot{\mathbf{a}}_m)$. These derivatives are evaluated from the coordinate time series using a five-point finite-difference stencil, accurate to fourth order in the sampling step (Brunton et al. 2016). Although finite differencing tends to amplify high-frequency noise, the truncation of the small-scale POD tail acts as an effective low-pass filter, rendering the retained coordinates sufficiently smooth that noise-robust differentiation schemes (Chartrand 2011) are not required. The resulting snapshot pairs $(\mathbf{a}_m, \dot{\mathbf{a}}_m)$ form the training dataset from which the continuous RBF vector field is constructed.

3 Reduced-order model of the wall shear

We model the temporal evolution of the wall shear stress in latent space using the reduced-order framework of Pérez Cuadrado et al. (2026), as illustrated in Fig. 2. For self-containment, the formulation is

summarised hereafter. A single global vector field advances the latent coordinates $\boldsymbol{a}$ (Section 2),

$$\dot{\boldsymbol{a}} = f(\boldsymbol{a}) = \sum_{i=1}^{n_c} \boldsymbol{\xi}_i \, \phi_i(\boldsymbol{a}), \tag{4}$$

expressed as a weighted sum of n_c anisotropic Gaussian radial basis functions (RBFs). Each function is centred at $\boldsymbol{c}_i$ and shaped by the local phase-space geometry of the attractor,

$$\phi_i(\boldsymbol{a}) = \exp\!\left(-\tfrac{1}{2}(\boldsymbol{a}-\boldsymbol{c}_i)^\top \mathbf{M}_{k(i)}(\boldsymbol{a}-\boldsymbol{c}_i)\right), \qquad \mathbf{M}_k = \mathbf{C}_k^{-1}, \tag{5}$$

where the trajectory (resampled for constant arc lengths) is partitioned into K clusters via K-means clustering. Here, $k(i) \in \{1, \ldots, K\}$ denotes the cluster assignment of RBF centre $\boldsymbol{c}_i$, and $\mathbf{C}_k$ is the covariance matrix of cluster k, which shapes the local anisotropic metric $\mathbf{M}_k = \mathbf{C}_k^{-1}$. The precision $\mathbf{M}_k = \mathbf{C}_k^{-1}$ elongates each kernel along the attractor's local tangent directions and narrows it along the normal ones, so that the union of the RBFs follows the manifold's shape and orientation. The weights $\boldsymbol{\xi}_i$ are fitted once, by a single global ridge regression over all snapshots, giving one continuous, differentiable field with no switching between local models. The kernel widths and the placement of the centres within each cluster are set out in Section 4, and further construction details can be found in Pérez Cuadrado et al. (2026).

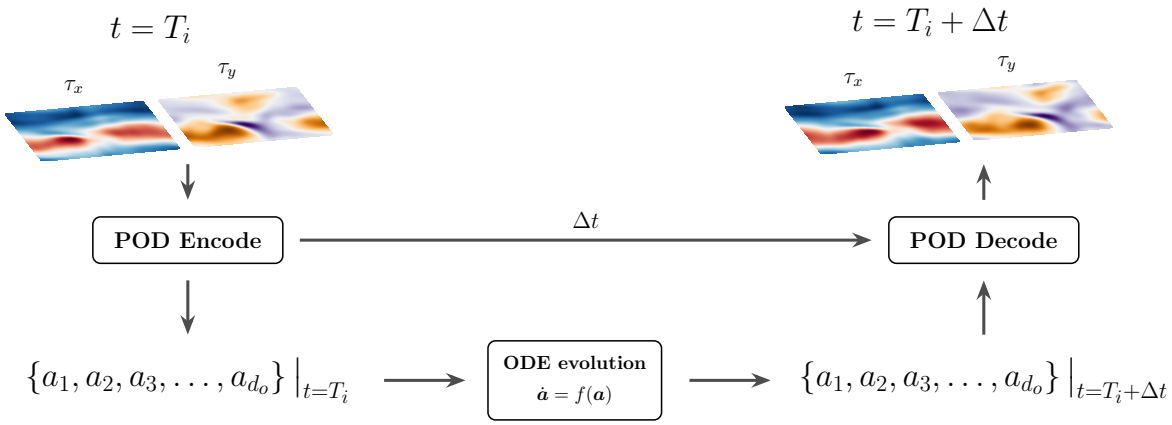


Fig. 2 Modelling pipeline. The wall shear-stress components (τ_x, τ_y) at time T_i are encoded onto latent coordinates by proper orthogonal decomposition, advanced by the learned vector field $\dot{\boldsymbol{a}} = f(\boldsymbol{a})$, and decoded back to the wall shear-stress fluctuation at $T_i + \Delta t$

To identify distinct dynamical states on the attractor, the latent trajectory is partitioned using K-means clustering (Lloyd 1982). The optimal number of clusters K is determined by evaluating the separation between neighbouring phase-space regions using three diagnostic metrics. The primary spatial metric is the K-means Bayesian Information Criterion (Schwarz 1978; Pelleg and Moore 2000)

$$\mathrm{BIC}(K) = M \log J_K + K \log M - \frac{2}{d_o} \sum_{k=1}^{K} n_k \log \frac{n_k}{M}, \qquad J_K = \sum_{m=1}^{M} \left\| \mathbf{a}_m - \boldsymbol{\mu}_{k(m)} \right\|^2, \tag{6}$$

evaluated over M snapshots, where n_k and $\boldsymbol{\mu}_k$ denote the population size and mean of cluster k, respectively. The within-cluster variance J_K decreases as clusters are added, so $M \log J_K$ rewards partition tightness while $K \log M$ penalises model complexity. The final term is the cluster occupancy entropy; its $2/d_o$ prefactor minimises its influence, so it does not artificially shift the selection threshold. Rather than seeking a global minimum, the optimal K is selected at the prominent elbow of the discrete derivative $\Delta\mathrm{BIC}/\Delta K$ (Fig. 3a) (Satopaa et al. 2011).

The remaining two diagnostics evaluate the geometric properties of neighbouring clusters. The spatial elongation of cluster k is quantified by the participation ratio (Bell and Dean 1970) of its covariance

matrix $\mathbf{C}_k$,

$$P_k = \frac{\left(\operatorname{Tr}\mathbf{C}_k\right)^2}{\operatorname{Tr}\left(\mathbf{C}_k^2\right)} = \frac{\left(\sum_{p=1}^{d_o}\lambda_{k,p}\right)^2}{\sum_{p=1}^{d_o}\lambda_{k,p}^2}, \tag{7}$$

where $\lambda_{k,p}$ $(p = 1, \ldots, d_o)$ are the eigenvalues of $\mathbf{C}_k$. The participation ratio measures the effective number of directions across which the cluster spreads, ranging from $P_k \approx 1$ for a thin, elongated filament to $P_k \approx d_o$ for an isotropic cloud. Discrepancies in shape between adjacent clusters are tracked via the average elongation difference $\langle |P_k - P_{k'}| \rangle$ over neighbouring pairs, as shown in Fig. 3(b).

The local phase-space dynamics within each cluster are characterised by the mean trajectory flow speed,

$$v_k = \frac{1}{n_k} \sum_{m\,:\,k(m)=k} \|\dot{\mathbf{a}}_m\|, \tag{8}$$

which averages the trajectory velocity across all snapshots residing in cluster k. Neighbouring clusters k and k' exhibit distinct dynamical regimes when the normalised speed contrast,

$$\frac{|v_k - v_{k'}|}{v_k + v_{k'}} \in [0, 1], \tag{9}$$

is appreciable, separating quasi-stationary, high-residence regions from rapid, bursting transitions (Fig. 3c).

As the number of clusters K increases, all three metrics exhibit a sharp initial variation followed by a clear plateau. Beyond $K = 16$, the curves level off across all three state-space dimensions ($d_o = 45, 75, 170$) and oscillate within a narrow band. Although the asymptotic values depend on d_o, the qualitative shape of the curves and the location of the elbow remain identical. This invariance confirms that the $d_o = 45$ truncation already captures the core dynamical skeleton of the system.

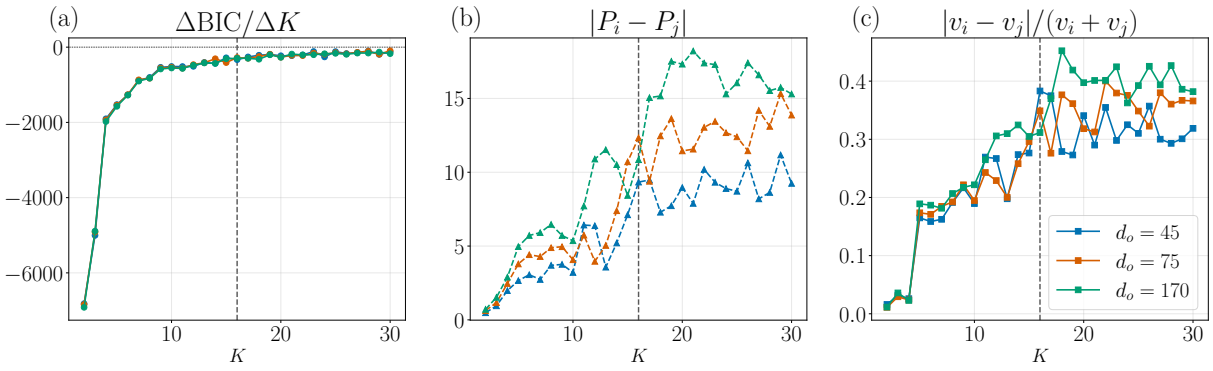


Fig. 3 Selecting the number of clusters. Three separation diagnostics against the number of clusters K, each for $d_o = 45$, 75 and 170, with the chosen $K = 16$ marked by the dashed line: (a) the marginal Bayesian information criterion $\Delta\mathrm{BIC}/\Delta K$; (b) the participation-ratio gap $|P_k - P_{k'}|$ between neighbouring clusters; (c) the relative speed gap $|v_k - v_{k'}|/(v_k + v_{k'})$

While these diagnostics help in determining the optimal number of clusters (i.e. sixteen, in the present case), they say nothing about their physical significance. Whether these clusters correspond to distinct stages of the near-wall cycle or are just arbitrary partition of the attractor depends on the flow physics: if each phase of the cycle is a genuine dynamical event, it must project a distinct scalar footprint onto the wall, reflecting the presence of streaks or quasi-streamwise vortices.

Streaks manifest as spanwise-alternating modulations of the streamwise velocity (Robinson 1991), making the streamwise shear component τ_x their primary footprint. Conversely, quasi-streamwise vortices imprint directly on the spanwise shear, reflecting the wall value of streamwise vorticity, $\tau_y = -\mu\omega_x|_{z=0}$. Following Hamilton et al. (1995), these two physical mechanisms are quantified by the scalar markers

$$A_x = \frac{1}{\tau_w}\left|\widehat{\langle\tau_x'\rangle_x}\,(k_y{=}1)\right|, \tag{10}$$

$$\sigma_y = \frac{1}{\tau_w}\left\langle \tau_y'^{\,2}\right\rangle^{1/2}, \tag{11}$$

where τ_w is the mean wall shear stress, $\langle\cdot\rangle_x$ denotes the streamwise average (filtering out $k_x \neq 0$), and $\widehat{(\cdot)}$ is the spanwise Fourier transform. Thus, A_x measures the fundamental ($k_y = 1$) spanwise amplitude of the streak, while σ_y is the root-mean-square of the spanwise shear fluctuations. On the streak–vortex phase plane (A_x, σ_y), the cyclic alternation of the near-wall dynamics should manifest as a closed orbit, with each cluster identifying a distinct phase of the cycle.

These cluster-averaged markers (standardised to zero mean and unit variance) are shown in Fig. 4(a), together with the dominant transitions between clusters. The occupied quadrants identify three phases of the self-sustaining regeneration cycle (Hamilton et al. 1995; Waleffe 1997): a quiescent state in Q3 (low A_x, low σ_y), a mature high-amplitude streak in Q1 (high A_x, moderate σ_y), and a vortex-dominated state in Q2 (high σ_y, low A_x).

This physical picture is consistent with the cluster-averaged wall-shear fields τ_x shown in Fig. 4(d–f). In Q1, the field exhibits a strong, streamwise-coherent low-speed streak. In Q2, the streak structure lacks streamwise uniformity, appearing instead as streamwise-modulated, localised patches enclosed by closed τ_x contours. In Q3, only a weak, low-amplitude remnant of the streak remains.

Even without physical labels during clustering, the resulting centroids coincide with the three phases of the near-wall cycle when evaluated through the wall markers. This separation stems from the manifold-adapted clustering of Pérez Cuadrado et al. (2026), which corrects for coordinate representation and trajectory sampling biases that would otherwise blur the phases.

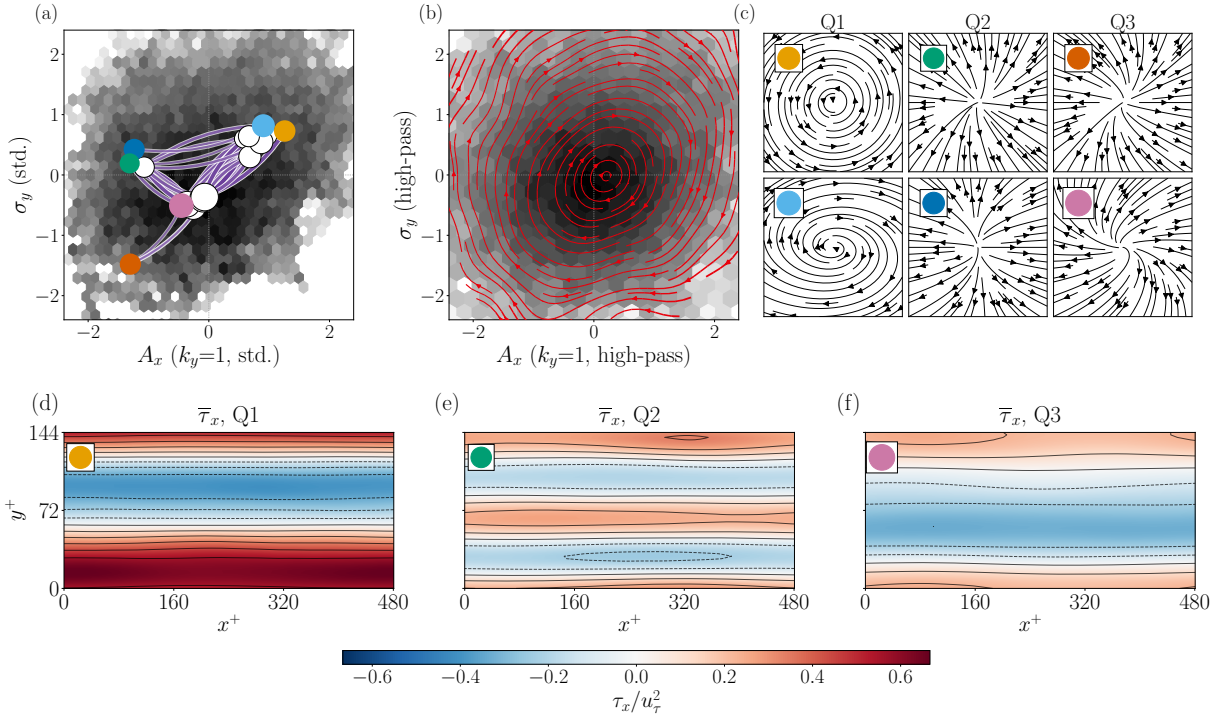


Fig. 4 The near-wall cycle on the streak–vortex plane. (a,b) Plane spanned by the coherent streak amplitude A_x, the $k_y = 1$ spanwise fundamental of τ_x', and the spanwise-shear intensity σ_y. (a) The sixteen cluster centroids on the standardised plane, sized by occupancy, over the sampled density (grey), with arrows for the dominant inter-cluster transitions. (b) The same plane after high-pass filtering, with red streamlines of the mean phase velocity tracing a weak clockwise circulation. (c) Local dynamics of the fitted field, the leading-plane projection of the Jacobian $\mathbf{J}_k$ at two clusters per quadrant, shown as black streamlines, with a corner marker colouring each tile to match its cluster in (a). (d–f) Cluster-averaged streamwise shear $\overline{\tau}_x$ of the representative cluster in Q1, Q2 and Q3: the mature streak, the vortex-dominated phase, and the quiescent state

The transitions between phases exhibit a clear directional sense, rendered in Fig. 4(b) by high-pass filtering both markers and plotting their mean phase-space velocity field as streamlines. The spanwise shear leads the streak amplitude by approximately $13\,t^+$ (roughly one-tenth of the cycle period), producing a clock-wise circulation in the phase plane. Thus, the cycle progresses sequentially from the quiescent state (Q3) through the vortex-dominated phase (Q2) to the mature streak (Q1), before returning to the quiescent state via a vortex-depleted, decaying-streak region in Q4. The rise of vortices prior to the development of the coherent streak is directly consistent with the lift-up mechanism (Brandt 2014).

The circulation is kinematic rather than causal: along a closed phase-space orbit, a phase lead merely reflects the sequential ordering of states, not the direction of dynamical forcing. Consequently, the observed phase sequence does not imply that lift-up actively drives the cycle. In fact, Lozano-Durán et al. (2020) demonstrated that the dominant causal drive operates in reverse, with streak breakdown actively regenerating the vortices.

While the clusters correspond to the phases of the cycle, spatial partitioning alone does not guarantee dynamical distinctness. A static discretisation may split a continuous attractor into contiguous regions that share identical local dynamics. Clustering is physically justified only when the local flow field, and not merely state-space location, distinguishes the phases (Pérez Cuadrado et al. 2026). Since the fitted vector field is differentiable, this distinction can be evaluated directly at each centroid through the local Jacobian matrix $\mathbf{J}_k$, which governs the deformation of passing trajectories.

At every cluster centroid, the leading eigenvalue of $\mathbf{J}_k$ has a positive real part, indicating a locally unstable direction that stretches nearby phase-space trajectories. The nature of this leading eigenpair clearly separates the phases (Fig. 4c): the mature streak states in Q1 are dominated by complex-conjugate pairs, corresponding to local phase-space rotation. The Q2 clusters are predominantly real-dominant, imparting pure strain without rotation. Q3, on the other hand, occupies an intermediate regime. This structural shift in the spectrum confirms that the clusters isolate dynamically distinct stages of the cycle. Consequently, while standard cluster-based reduced-order models represent phase dynamics purely via inter-cluster transition probabilities (Kaiser et al. 2014; Fernex et al. 2021), the present differentiable field exposes each phase as a distinct local deformation of the attractor.

4 Model training and autonomous prediction

Following the clustering process, the radial basis functions are adapted to the underlying structure of the attractor (Pérez Cuadrado et al. 2026). A single global regression over the RBFs then determines the right-hand-side vector field $\mathbf{f}(\mathbf{a})$ governing the dynamical system $\dot{\mathbf{a}} = \mathbf{f}(\mathbf{a})$. The first 80% of the time record is used for the regression, while the remaining 20% is held out for testing.

Each kernel inherits its orientation from the cluster at which it is placed. The tangential directions are the leading eigenvectors of the local covariance $\mathbf{C}_k$ that carry 99% of the intra-cluster variance; the remaining directions form the normal complement. The clustering of Section 3 therefore fixes these directions for every RBF in the cluster, but not the width along them. Along the tangent directions the per-direction scaling $\sqrt{\lambda_{k,p}}$ is retained, and the width is

$$\sigma_{k,p} = \alpha_{\tan}\sqrt{\lambda_{k,p}}, \tag{12}$$

where the multiplier $\alpha_{\tan}$ is common to every cluster and controls the overlap between neighbouring kernels. Along the normal directions the spread vanishes, and the width is instead the median separation between neighbouring centres projected on each direction, floored at a small fraction of the smallest tangent width (Pérez Cuadrado et al. 2026). The centres themselves are placed by farthest-point sampling (Gonzalez 1985) within each cluster until its quota is filled, independently of the metric so that the near-degenerate normal directions cannot distort the cover.

Two parameters of the library consequently remain free: the number of basis functions N and the multiplier $\alpha_{\tan}$. Their effect on the fit is measured on the test segment by the mode-averaged normalised error,

$$\mathrm{NRMSE} = \frac{1}{d_o}\sum_{i=1}^{d_o}\frac{\mathrm{rms}\big(\dot{a}_{i,\mathrm{ROM}} - \dot{a}_{i,\mathrm{true}}\big)}{\sigma_{\dot{a}_i}}, \tag{13}$$

where $\sigma_{\dot{a}_i}$ is the standard deviation of $\dot{a}_i$ over the same segment. We sweep both parameters for two truncations, $d_o = 45$ and 75, over the ranges given by the colour scales in Fig. 5(a,b). The number of basis functions spans the same range in both cases, whereas $\alpha_{\tan}$ reaches higher values at $d_o = 75$, since a higher-dimensional latent space is more sparsely sampled and therefore requires broader kernels. The 99%-energy truncation $d_o = 170$ adds only the incoherent small scales already identified in Section 2; this truncation is retained in the clustering diagnostics of Section 3 but not carried into the dynamical model.

At small $\alpha_{\tan}$, the kernels are narrower than the spacing between their centres, and the envelope of RBFs leaves gaps across the manifold. Coverage then improves only slowly with N, whereas for the smallest widths no number of kernels within the swept range fills the space. As $\alpha_{\tan}$ grows, the supports overlap, the manifold is covered, and the error decreases. Once the width is sufficient, each curve reaches a minimum at an intermediate N and rises beyond it, the fit passing from under-parameterised to over-parameterised as the extra kernels begin to fit noise. The larger the width, the fewer kernels are needed to reach that point.

The minimum error of each curve across N is shown against $\alpha_{\tan}$ in Fig. 5(c), where the two truncation levels intersect. At small $\alpha_{\tan}$, the lower-dimensional model reaches a lower minimum because, with fewer modes, the spread of characteristic scales across dimensions is narrower. Thus, a single global $\alpha_{\tan}$ fits every dimension well. Past $\alpha_{\tan} \approx 10$, where coverage is no longer the limiting constraint, the higher-dimensional model achieves lower error because retaining more modes leaves less energy in the truncated tail, and a smaller truncation loss sets the error limit.

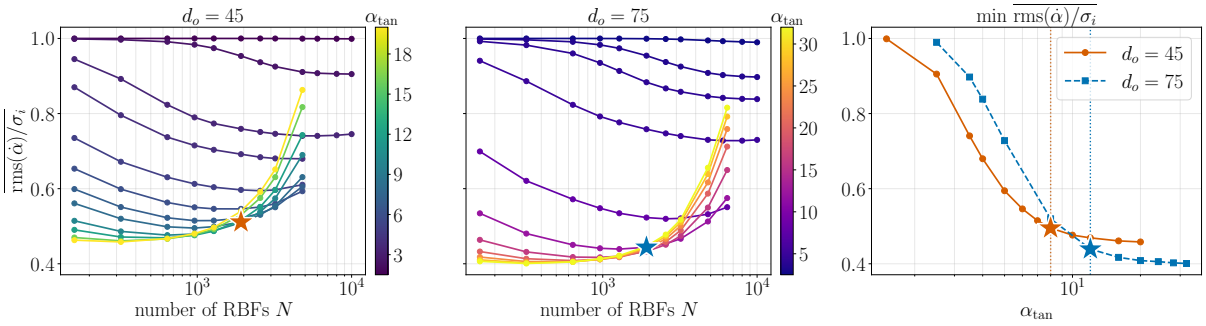


Fig. 5 Fitting the vector field. Mode-averaged test NRMSE against the number of radial basis functions N, one curve per tangent anisotropy $\alpha_{\tan}$ (colour), for (a) $d_o = 45$ and (b) $d_o = 75$. (c) The NRMSE floor over N against $\alpha_{\tan}$. Stars mark the selected models, $\alpha_{\tan} = 8$ for $d_o = 45$ and $\alpha_{\tan} = 12$ for $d_o = 75$, both at $N = 1920$

With the RBF parameterisation selected, the continuous vector field $\dot{a} = f(a)$ is fully defined and ready for time integration. Predictions are computed using a fourth-order Runge–Kutta scheme at the sampling interval $\Delta t \approx 0.53\, t^+$. To prevent unphysical trajectory drift in regions where state-space data are sparse, the integration is coupled with the per-cluster trust-region corrector of Pérez Cuadrado et al. (2026), calibrated as described therein. Each cluster trust region is defined to enclose the sampled dataset; consequently, trajectories remaining within the populated state space never trigger the corrector. The corrector functions solely as a safeguard against excursions beyond the sampled state space, pulling rogue trajectories back to keep the integration bounded. Across all unforced autonomous rollouts reported here, the corrector was never activated. The presented statistics are therefore those of the unconstrained fitted field itself.

Autonomous predictions are evaluated using two indicators. The first measures short-range accuracy through the valid prediction time (VPT), a standard gauge of how long a forecast remains close to the reference trajectory (Vlachas et al. 2020). It is defined as the time in t^+ before the reconstruction error e_{rec} first exceeds 0.25,

$$\mathrm{VPT} = \min\{\, t : e_{\mathrm{rec}}(t) > 0.25 \,\}, \qquad e_{\mathrm{rec}}(t) = \frac{\|\boldsymbol{a}_{\mathrm{ROM}}(t) - \boldsymbol{a}_{\mathrm{true}}(t)\|}{\sigma_A}. \tag{14}$$

Here e_{rec} is the latent-space error normalised by $\sigma_A = \big(\sum_i \mathrm{Var}(a_i)\big)^{1/2}$, the per-mode standard deviations over the reference record summed in quadrature.

The normalisation is global, by the single σ_A rather than per mode, so the error is dominated by the energetic modes that carry the wall signature, their absolute deviations being largest.

The quality of the long-time statistics is assessed using two metrics that compare the autonomous rollout against the reference trajectory (Gibbs and Su 2002). Defining $s(t) = \langle \tau_x \rangle(t)$ as the instantaneous mean skin friction and F as its empirical cumulative distribution over the predicted time series, the first metric is the Wasserstein-1 distance (Peyré and Cuturi 2019),

$$W_1 = \int \big| F_{\mathrm{ROM}}(s) - F_{\mathrm{FOM}}(s) \big| \, \mathrm{d}s, \tag{15}$$

comparing the reduced-order model (ROM) against the full-order model (FOM) reference, reported normalised by the reference standard deviation σ_{FOM}. The second metric compares the shape of the skin-friction power spectrum. Let $\widehat{\phi}(f)$ be the premultiplied power spectral density $f\,P(f)$ of each series, normalised to unit area over the resolved band $f \le f_{\max}$, above which the reference spectrum is band-limited. The two spectral shapes are compared using the total-variation distance between them, treating

the distributions as probability densities over frequency (Gibbs and Su 2002),

$$\mathrm{TV} = \tfrac{1}{2} \int \left| \widehat{\phi}_{\mathrm{ROM}}(f) - \widehat{\phi}_{\mathrm{FOM}}(f) \right| \mathrm{d}f. \tag{16}$$

The factor of 1/2 together with the unit-area normalisation bounds TV to the interval $[0, 1]$, yielding zero for identical spectral shapes and one for disjoint support. Since it is area-normalised, TV measures the spectral shape alone, complementing W_1, which quantifies discrepancies in the overall amplitude distribution of skin friction. Both spectra are estimated using Welch's method with identical windowing and segment overlap on a common frequency grid.

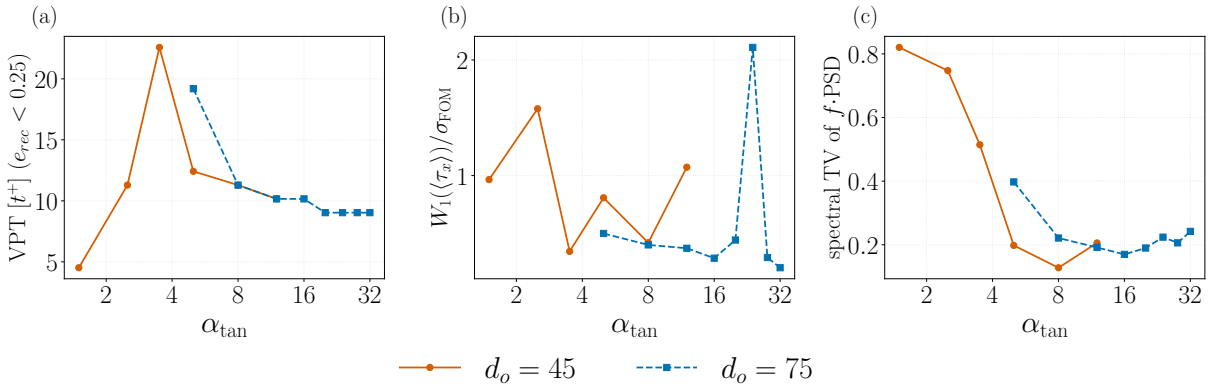


Fig. 6 Predictability and climate against anisotropy. For $d_o = 45$ and 75 against $\alpha_{\tan}$: (a) valid prediction time to reconstruction error $e_{\mathrm{rec}} = 0.25$; (b) Wasserstein-1 distance of the skin-friction distribution in units of σ_{FOM}; (c) total-variation distance of the premultiplied spectrum $f \cdot \mathrm{PSD}$

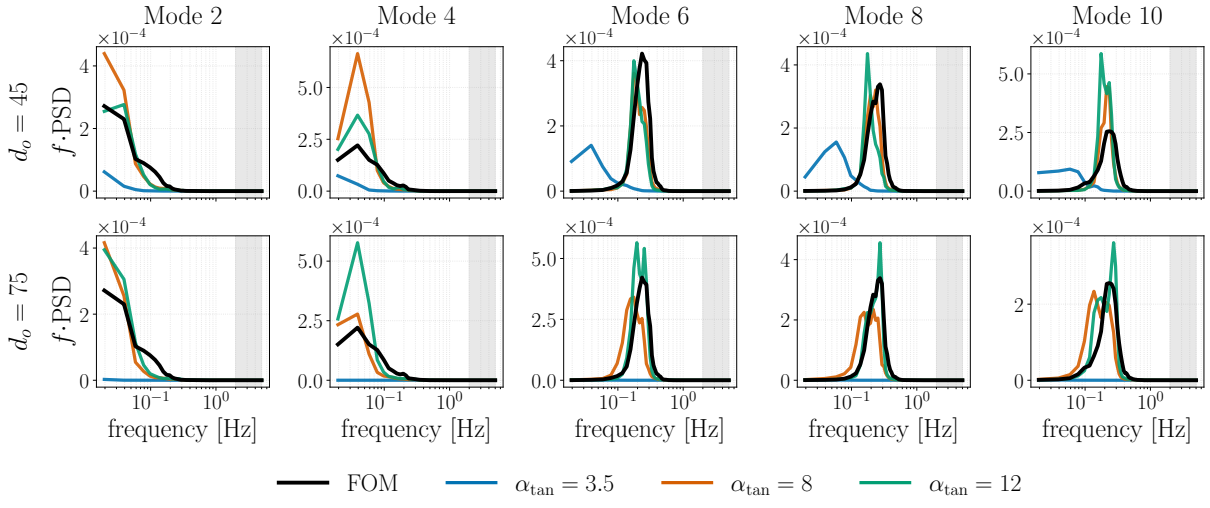


Fig. 7 Spectral fidelity of the free run. Premultiplied spectra $f \cdot \mathrm{PSD}$ of latent modes 2, 4, 6, 8 and 10 for the free-integrated model at three values of $\alpha_{\tan}$ and the reference (FOM), for $d_o = 45$ (top) and 75 (bottom). Shading marks frequencies above the reference band limit

Fig. 6 evaluates model performance across three indicators, i.e. the valid prediction time (VPT), the skin-friction distribution error (W_1), and the spectral shape distance (TV), plotted against $\alpha_{\tan}$ for both values of d_o. The valid prediction time, averaged over 100 initial conditions within the test segment, is presented in Fig. 6(a). Owing to the distinction between one-step fitting accuracy and closed-loop rollout performance, the VPT trend does not track the normalised root-mean-square error (NRMSE) minimum observed in Fig. 5(c). While the single-step fit improves monotonically with $\alpha_{\tan}$, the VPT peaks at an intermediate kernel width ($\alpha_{\tan} \approx 3.5$–5) and degrades thereafter for both values of d_o. This optimal

width remains unchanged when varying the error threshold between 0.2 and 0.35. The observed ranking therefore does not depend sensitively on the chosen tolerance.

Similarly, the long-term statistical state does not track the VPT. A larger W_1 yields a greater deviation from the reference skin-friction distribution, with its minimum occurring at a larger $\alpha_{\tan}$ than the VPT peak (panel b). The spectral distance exhibits the same trend (panel c). Since the one-step fit, predictability, and long-time statistics reach optimal values at different widths, selecting a single global $\alpha_{\tan}$ requires a compromise among these competing criteria.

This deficiency arises directly from the planar averaging operation itself. Since it retains only the spatial-mean component of each mode, the planar-averaged skin friction $\langle \tau_x \rangle$ is sustained by a small subset of low-order modes. The energetic, spanwise-alternating streak pairs average to near-zero across the span and yield a negligible contribution. At small values of $\alpha_{\tan}$, coverage gaps render the fitted field incapable of driving these modes, causing the unforced evolution to relax toward a quasi-steady state and resulting in a weak $\langle \tau_x \rangle$ (Fig. 8). As the width increases, the modes regain dynamic activity; however, within $\alpha_{\tan} \approx 3.5$–5, the active low-order modes oscillate at lower frequencies than the reference (Fig. 7). Consequently, the reconstructed $\langle \tau_x \rangle$ exhibits a shift toward lower frequencies rather than a simple attenuation of high-frequency content.

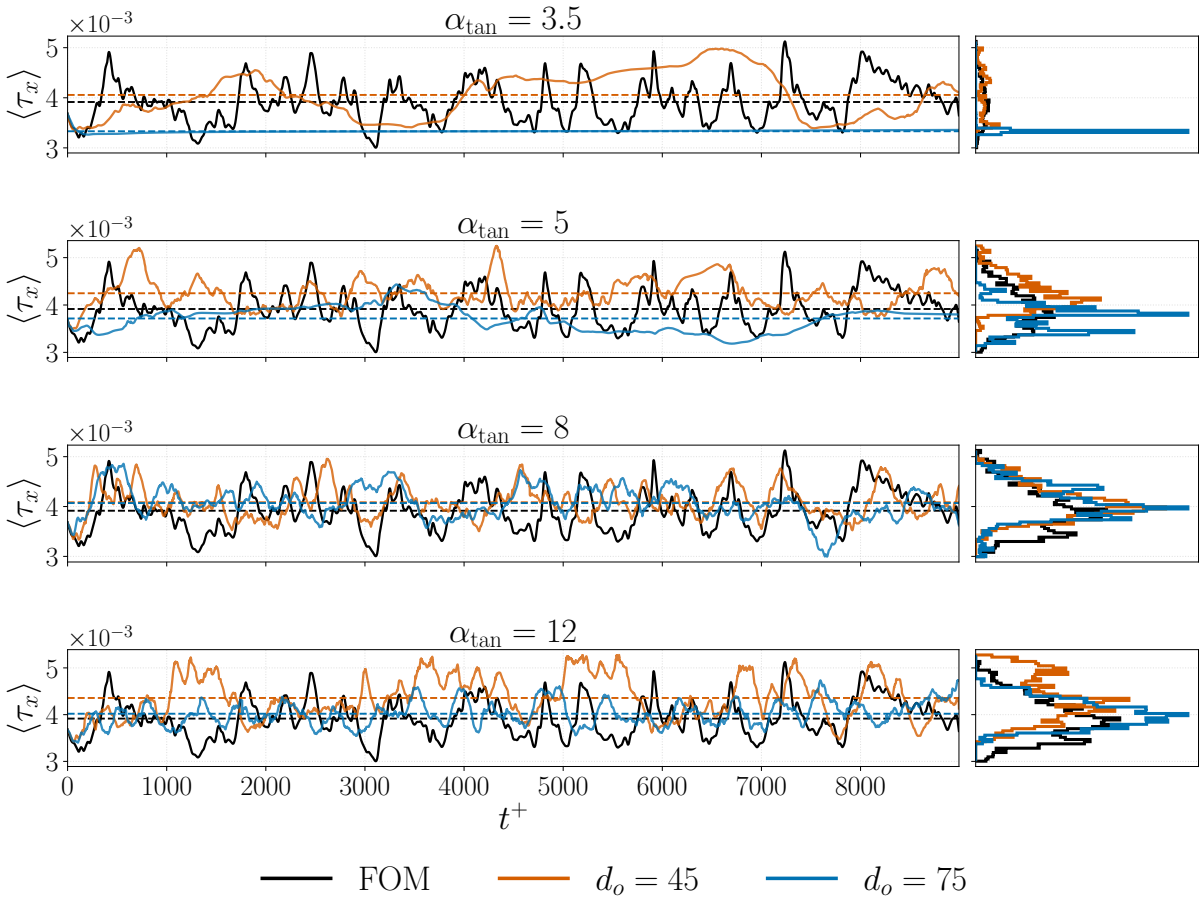


Fig. 8 Free-running skin friction. Spatially averaged streamwise wall shear $\langle \tau_x \rangle$ against time for the free-integrated model at $d_o = 45$ and 75 and the reference (FOM), one row per $\alpha_{\tan}$, with the time means dashed. Right: the marginal distributions

In contrast, the long-term statistical state improves at larger kernel widths, with the skin-friction and spectral distances reaching their minima around $\alpha_{\tan} \approx 8$–16. Wider, more strongly overlapping kernels regularise the fitted field; this sacrifices short-term spatial precision and predictability to ensure a bounded, accurate asymptotic state. This effect mirrors the compromise between short-term accuracy and long-time stability encountered in dissipated POD–Galerkin reductions (Aubry et al. 1988; Rempfer 2000), though here it emerges directly from the spatial kernel width rather than an explicit damping term.

Increasing the modal truncation order does not change this outcome. Across all three evaluation metrics, expanding the system dimension from $d_o = 45$ to $d_o = 75$ yields no performance improvement. The additional modes introduced in this range remain dynamically inactive with respect to the wall-shear

statistics. This observation justifies omitting higher-order modes up to $d_o = 170$, since incorporating those fine-scale components improves instantaneous spatial reconstruction without enriching the underlying physical dynamics.

In summary, narrow kernels maximise short-term trajectory tracking and prediction horizons, whereas wide kernels regularise the fitted field to correctly capture long-term statistical distributions.

These competing requirements can be balanced to design the operating point. Since a wall-only model is intended to recover the invariant measure of the near-wall cycle rather than individual trajectories, the kernel width is selected to optimise long-time statistics. Accordingly, the model is configured at $d_o = 45$ and $\alpha_{\tan} = 8$, representing the lower limit of the regime where the skin-friction distribution and spectral properties match the reference data most closely. This choice keeps performance near the predictability peak, albeit with a shorter prediction horizon than the width that maximises trajectory accuracy. At this parameter setting, the free integration accurately reproduces the skin-friction distribution and pre-multiplied power spectrum of the direct numerical simulation (Figs. 8 and 7). The wall-shear field, when integrated independently, sustains the statistical features of the near-wall cycle.

5 Dynamical analysis of the fitted field

The differentiability of the fitted field allows for a direct dynamical systems analysis beyond simple forward simulation. Resting states, system Jacobians, and perturbation growth can be evaluated in closed form. This section explores these features to map the boundaries of information retrievable from wall measurements alone.

Since the field $\dot{\boldsymbol{a}} = f(\boldsymbol{a})$ is defined over the entire state space, its slow points (i.e. the fixed points) can be located directly. A slow point is a state where the model tendency vanishes, corresponding to an equilibrium of the fitted field near which the flow evolves slowly. Since the local dynamics associated with these equilibria exhibit both stable and unstable directions, they act as saddle points that constitute weakly unstable near-equilibria.

Solving $f(\boldsymbol{a}) = 0$ for the production model reveals that every recovered slow point resides in the quiescent half of the cycle ($\sigma_y < 0$) and is distributed across a range of streak amplitudes. This arrangement is reminiscent of exact coherent structures in wall-bounded shear flows (Nagata 1990; Waleffe 2001; Toh and Itano 2003), as illustrated in Fig. 9a. Since these roots lie well within the region of densely sampled data where the corrector term remains inactive, they are not artefacts of sparse kernel coverage.

This arrangement persists across four independently fitted models, whose escape-time distributions show close agreement (Fig. 9b). The temporal rhythm of the cycle is inherently asymmetric: the quiescent phase provides weakly unstable anchors where trajectories linger for tens of viscous time units (t^+), whereas the active phase is devoid of such structures and is traversed rapidly. Unstable invariant solutions of the full near-wall cycle exhibit a similar dynamical rhythm, appearing as shadowed states interleaved between rapid bursts (Kawahara and Kida 2001; Kawahara et al. 2012). However, these slow points of the wall-visible projection are neither exact equilibria of the governing fluid equations nor direct representations of the true full-state invariant solutions.

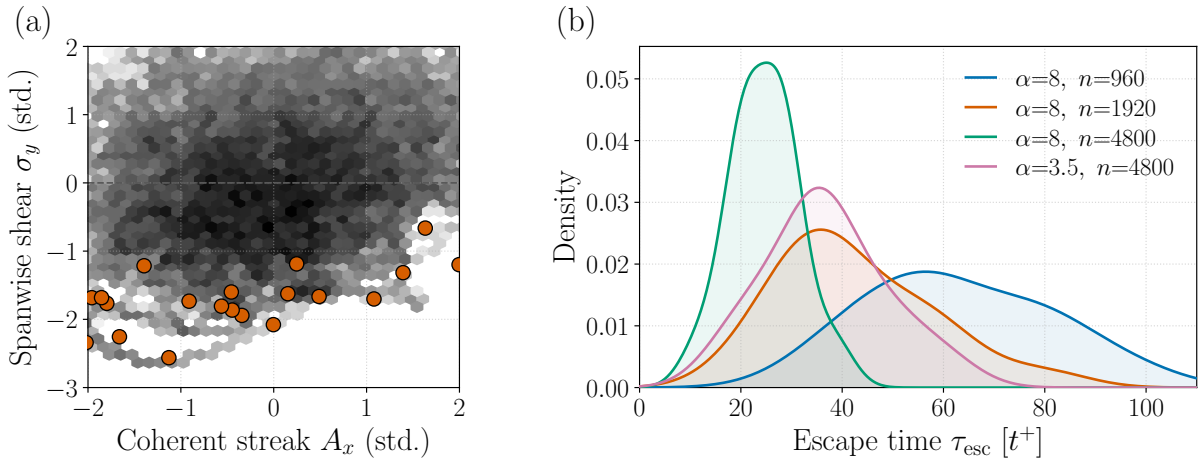


Fig. 9 Slow points live only in the quiescent phase. (a) States at which the fitted field's mean tendency vanishes, for the production model ($\alpha_{\tan} = 8$, 1920 kernels), on the streak–vortex plane over the sampled density (grey). All lie in the quiescent half ($\sigma_y < 0$), the active half none. (b) Escape-time distribution of the slow points for four independently fitted models, varying $\alpha_{\tan}$ and kernel count n: tens of t^+ in every case

The equilibrium slow points determine the regions where the flow velocity tends toward zero, whereas the Jacobian describes the local stretching and squeezing of adjacent trajectories. Differentiating the field yields the Jacobian $J = \partial f/\partial \boldsymbol{a}$, and its linearisation reveals the local structure of the cycle (Fig. 10). Beyond the eigenvalue distribution examined previously, the non-normality of J quantifies the worst-case transient amplification of disturbances over a finite horizon, while its trace indicates whether a local ensemble of neighbouring states contracts or expands.

Non-normal growth characterises the sensitivity of the flow to perturbations. Over a $10\,t^+$ horizon, G is the maximum factor by which a small perturbation amplifies in the latent-coefficient energy norm, defined as the leading singular value of the tangent-linear propagator integrated along the trajectory, not at a frozen point (Trefethen et al. 1993; Schmid 2007). Formulated in this manner, it captures transient growth permitted by the non-normal operator even when the real parts of the eigenvalues are negative (Butler and Farrell 1992). A large value of G identifies a sensitive phase where minor disturbances experience rapid growth. When evaluated across clusters (Fig. 10a), G peaks at the onset of the burst, where a disturbance amplifies by a factor of approximately 2.5. Since this growth is measured strictly in the wall-visible coordinates using a latent norm, G is a wall-projected proxy for physical energy growth rather than an absolute bound.

The trace of J characterises the rate of volume expansion or contraction for an ensemble of states rather than an isolated perturbation. As shown in Fig. 10a, phase-space squeezing is strongest during the mature burst phase and diminishes nearly to zero at burst onset. Since the sign of $\mathrm{tr}\,J$ exhibits minor variations across independent fits, only the relative ordering is interpreted; this metric also reflects the contraction of the wall-visible latent volume, not the true volumetric dissipation of the fluid.

When mapped onto the cycle, the most amplifying phase corresponds to the mature streak (Q1): the streak is most susceptible to perturbations precisely when it reaches peak amplitude prior to breakdown. Conversely, the quiescent state (Q3) exhibits the lowest sensitivity. The strongest phase-space contraction occurs during the decay of the mature streak, where amplified disturbances dissipate as the system relaxes back toward the quiescent regime.

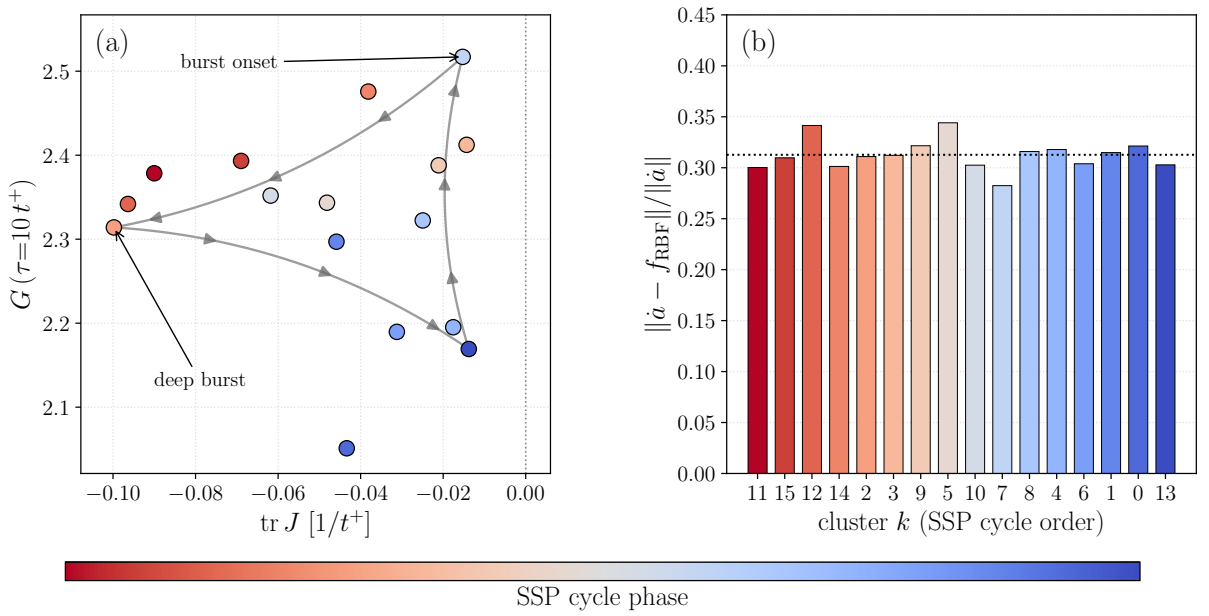


Fig. 10 Dynamical anatomy of the cycle, read off the fitted field. (a) Non-normal growth G, the worst-case amplification over a $10\,t^+$ horizon, against phase-volume contraction $\mathrm{tr}\,J$, one point per cluster, coloured by cycle phase. Burst onset is least contracting and most amplifying (upper right), the deep burst most contracting. (b) Closure deficit per cluster: the fraction of the wall tendency not determined by the wall state, near 0.3 and phase-uniform

The readings examined thus far reflect strictly what can be inferred from the wall-visible field. A fundamental limitation of this perspective is quantified by the unexplained residual shown in Fig. 10b: approximately 30% of the instantaneous wall tendency cannot be expressed as a function of the wall state alone, a proportion that remains uniform across all phases of the cycle. This uncaptured fraction is the portion of the future dynamics that the wall cannot discern from its current observations. It is driven primarily by unobserved interior motions, such as sweeping and ejections occurring away from the

boundary, and is an upper bound, as it also absorbs errors from POD truncation and finite-difference approximations. This hidden variability explains why the wall-visible model successfully reproduces the long-term statistics of the cycle yet fails to track individual trajectories deterministically. The subsequent section evaluates how this single-step deficit accumulates over the assimilation window.

6 Data assimilation and wall-observable autonomy

The unforced simulation of Section 4 reproduces the cycle's statistics, and with them its invariant measure. The near-wall region sustains its cycle when the motions above the buffer layer are quenched (Jiménez and Pinelli 1999), which justifies a low-dimensional model of it. Whether the wall shear is a sufficient coordinate for this cycle is a central question of the present contribution and, more generally, for the near-wall turbulence modelling community.

The cycle turns on the buffer-layer state above the wall: the wall-normal velocity that drives the lift-up, and the streamwise vortices that regenerate the streaks (Waleffe 1997; Hamilton et al. 1995). Since the wall carries only their footprint, a field fitted in wall coordinates learns a conditional-mean tendency (Chorin et al. 2000), averaged over the interior motions it does not see. To determine whether this wall-only representation is truly sufficient, we must look beyond unforced evolution. Section 5 read this averaging off as a closure deficit at a single step. What follows measures the forecast error it produces over short windows, quantifying the cost the wall-only model pays for the events off the wall, and whether online correction can overcome it.

A state estimator provides the natural framework to correct the deficit of a wall-observation-only model online. The wall shear is what an experiment can measure, even with sparse and noisy sensors, and the forecast is corrected online, without re-solving the model or forming an adjoint. Operating at a $5\,t^+$ assimilation cadence, the implementation runs $1.3\times$ faster than the physical flow it tracks, and is therefore practical for real-time deployment against live wall sensors. Assimilating wall measurements into a flow model has precedent with the governing equations in the loop (Colburn et al. 2011) and with a data-driven reduced model in their place (Özalp et al. 2026), the reduced model costing a fraction of the computational overhead required when the full equations are in the loop.

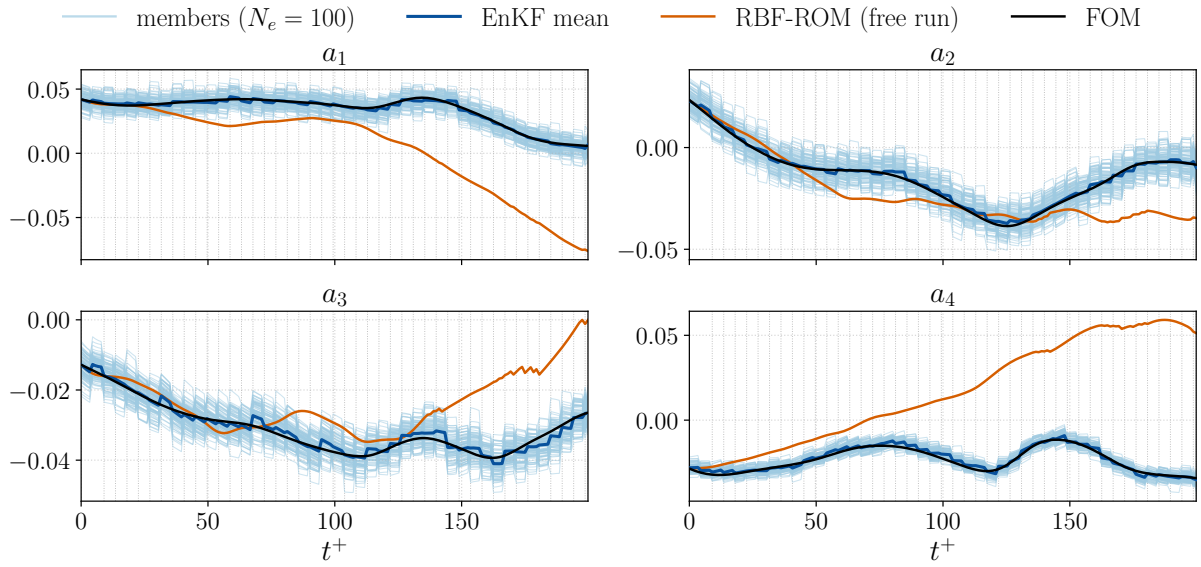


Fig. 11 Ensemble forecast of the leading modes. Latent coefficients a_1 to a_4 against t^+ for the ensemble Kalman filter ($N_e = 100$ members), its mean, the free-running RBF-ROM and the reference (FOM). Dotted verticals mark the analysis times

To implement this, we employ a Kalman filter as our estimator. It alternates a forecast (i.e. our ROM), in which the model advances the state and its uncertainty, with an analysis, in which the forecast is pulled towards the observations, the two weighted by their covariances through the Kalman gain (Kalman 1960). The classical formulation, however, propagates the full covariance through the linearised dynamics. To avoid linearisation, an ensemble Kalman filter (EnKF) carries a Monte Carlo ensemble, integrates each member through the nonlinear field, and reads the forecast covariance from their spread (Evensen 2003).

This captures the nonlinear error growth along the attractor, at a cost set by the ensemble size, the latent dimension and the number of kernels, $O(N_{\text{ens}}\, d_o\, n_c)$, and independent of the DNS grid. Since the computational cost of solving full DNS equations scales steeply with Reynolds number, while the cost of our reduced-order ensemble remains tied only to the latent space, the efficiency margin of our approach has the potential to widen substantially at higher Reynolds numbers.

To configure this filter for our channel flow configuration, the estimation operates directly within the reduced coordinate space. The filter runs on the reduced coordinates, observing the two wall-shear components on a uniform grid of stride Δ, so $N_{\text{obs}} = 2(64/\Delta)^2$ per sample. Each sensor carries independent Gaussian noise of standard deviation σ_{obs} relative to the root-mean-square of the wall-shear fluctuation, yielding a diagonal observation covariance. To ensure the ensemble members do not artificially converge, a unique realisation of this sensor noise is added to the observations for each individual model member (i.e., parallel runs of our ROM).

Table 1 Baseline data-assimilation parameters

N_{ens}	Δ	Components	N_{obs}	σ_{obs}	$\Delta t^{+}_{\text{obs}}$	Mult. infl.	q
100	8	τ_x, τ_y	128	0.1	5	1.02	0.1

An ensemble of N_{ens} members is initialised by perturbing a snapshot about the attractor, advanced through the fitted field between analyses, and corrected every $\Delta t^{+}_{\text{obs}}$. Multiplicative and additive inflation counter ensemble collapse, the additive level q being the control variable examined below. Table 1 lists the baseline values, with the stride, noise level and cadence varied later. On the leading modes (Fig. 11) the ensemble brackets the reference, and its mean tracks the true trajectory where the free run, given the same start, drifts away.

At each analysis step, the filter compares the incoming wall observations with its own forecast. The mismatch between the two—known as the innovation d (i.e., the difference between the observation and the ROM forecast)—is the primary quantity it weighs against its internal error estimate. To determine whether this mismatch exceeds the uncertainty the filter expected, d is scaled by the predicted innovation covariance $S = \mathbf{H}\mathbf{P}^f\mathbf{H}^\top + \mathbf{R}$, yielding the normalised innovation squared, defined as $\chi^2 = d^\top S^{-1} d$. Here, $\mathbf{H}$ is the linear map from the reduced state to the sensors, $\mathbf{P}^f$ denotes the forecast covariance derived from the ensemble spread, and $\mathbf{R}$ stands for the diagonal observation-noise covariance.

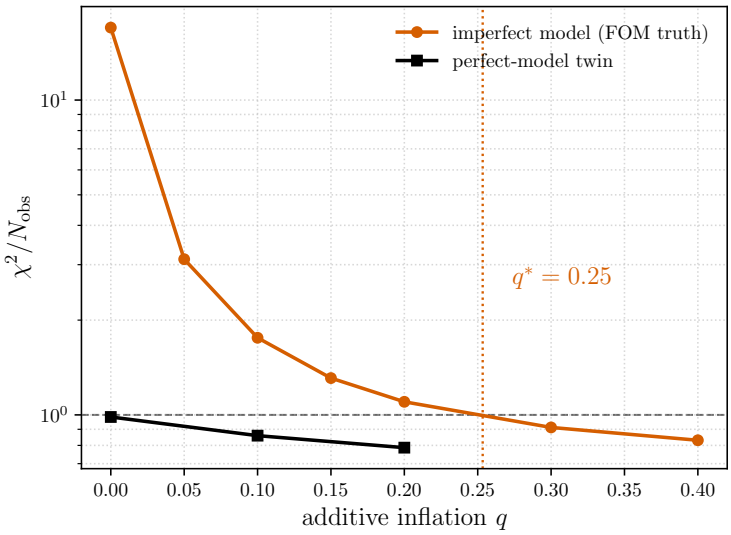


Fig. 12 The additive-inflation floor. Time-averaged χ^2/N_{obs} against the additive inflation q, for the model assimilating FOM observations and for the perfect-model twin. The crossing of $\chi^2/N_{\text{obs}} = 1$ (dashed) sets q^*. The twin needs none

When the filter is consistent, this normalised statistic averages to one per observation, $\chi^2/N_{\text{obs}} = 1$, which is the standard innovation-consistency check (Desroziers et al. 2005).

If the model advancing the ensemble misses certain flow dynamics, it produces innovations that are too large for the predicted spread, resulting in $\chi^2/N_{\text{obs}} > 1$. Enlarging the forecast covariance via additive

inflation q restores consistency. The specific level that brings χ^2/N_{obs} back to unity is denoted as q^*: this is the artificial process noise the wall-only model must inject to account for its own structural errors, establishing a standard approach for quantifying model deficiencies directly from the innovations (Dee 1995; Anderson 2007; Tandeo et al. 2020).

Building on this definition, the value of q^* carries a clear physical meaning. It is the forecast error that the wall-only model incurs due to the unresolved interior motions. This error is expressed in units of the attractor's climatological standard deviation, i.e. units of the attractor's long-term standard deviation. The inflation q^* is an upper bound rather than an exact equality, because it also absorbs errors from POD truncation and finite-difference derivatives. Since these auxiliary errors compound the discrepancy, the true contribution of the missing interior motions alone can be no larger than this limit.

As shown in Fig. 12, this value is located precisely at the crossing where $\chi^2/N_{\mathrm{obs}} = 1$. A perfect-model twin experiment, where the "truth" is generated by the model itself, requires no such inflation: its q^* is zero. The error floor measured against the actual flow is therefore genuine model deficiency rather than a filtering artefact. This twin framework also separates the two distinct reasons a wall-only forecast decays: chaotic trajectory divergence (which affects any nonlinear dynamical model) versus the missing interior information quantified by q^*. For the present system, $q^* = 0.25$, which contrasts with the operational $q = 0.1$ from Table 1 used to advance the ensemble. Consequently, over each $5\,t^+$ interval, the model must inject process noise with an amplitude equal to one-quarter of the climatological standard deviation (amounting to a few percent of the climatological variance) to maintain statistical consistency with the observations.

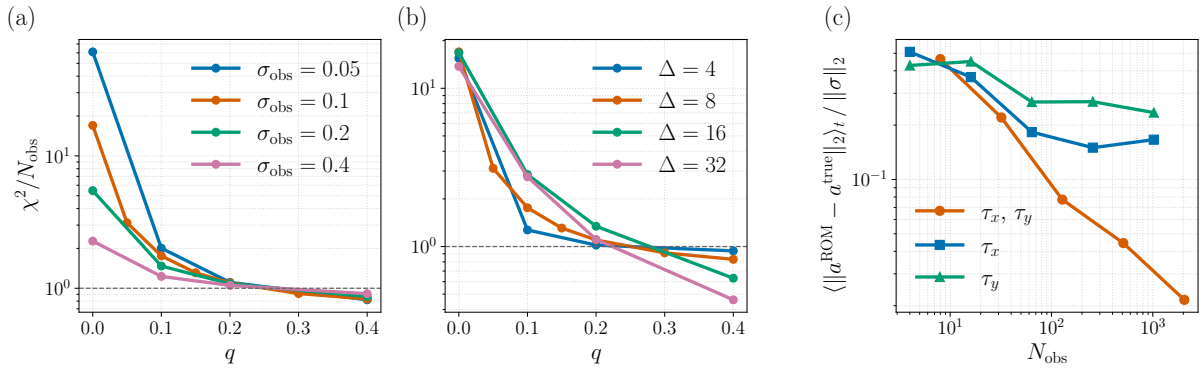


Fig. 13 q^* is a property of the model. χ^2/N_{obs} against q across (a) observation-noise level σ_{obs} and (b) sensor stride Δ. The unit crossing barely moves. (c) Analysis error, normalised by climatology, against sensor count N_{obs} for τ_x, τ_y and both components

To separate this intrinsic model error from the way the wall is observed, q^* is tested across different measurement configurations, including variations in both the observation-noise level and the sensor stride (Fig. 13). When either the observation-noise level (a) or the sensor stride (b) is varied over an eightfold range, the point where $\chi^2/N_{\mathrm{obs}} = 1$ barely moves, holding steady near $q^* \approx 0.25$.

Raising the observation noise only lowers the curve at $q = 0$ (because a larger covariance matrix $\mathbf{R}$ makes the same innovations less surprising to the filter), while leaving the critical crossing point completely unchanged. Changing the sensor stride is more drastic: at a coarse stride of $\Delta = 32$, the wall resolves only 8 of the 45 modes (noted as the rank of $\mathbf{H}$ in the legend), yet q^* remains invariant. This inherent mismatch is therefore a fundamental property of the model itself rather than an artefact of the filter or the sensor network, which is what allows it to be reliably interpreted as the model's own structural deficiency.

The state estimate itself does depend directly on the sensing setup. Panel (c) shows the analysis error, normalised by the long-term spread, decreasing as more sensors are added. Using both shear components together outperforms either component alone, and τ_x carries more information than τ_y, because the fluid streaks leave a stronger streamwise footprint than the quasi-streamwise vortices leave in the spanwise shear. Denser, two-component sensing thus sharpens the state estimate without altering this baseline error: how well the wall is read and how much the wall omits are entirely separate quantities.

The inflation q^* is a time average, and the deficit it measures is not spread evenly over the cycle. Resolved in time over a $2000\,t^+$ run, the excess above consistency arrives in bursts rather than as a steady offset (Fig. 14a). Averaged per cluster on the streak–vortex plane, it is largest in the active, high-σ_y phase and near-consistent in the quiescent phase (Fig. 14b). The wall-only forecast therefore stays consistent

with the observations through the quiescent phase and falls short in the active one, where the lift-up mechanism and streak breakdown are governed from off the wall.

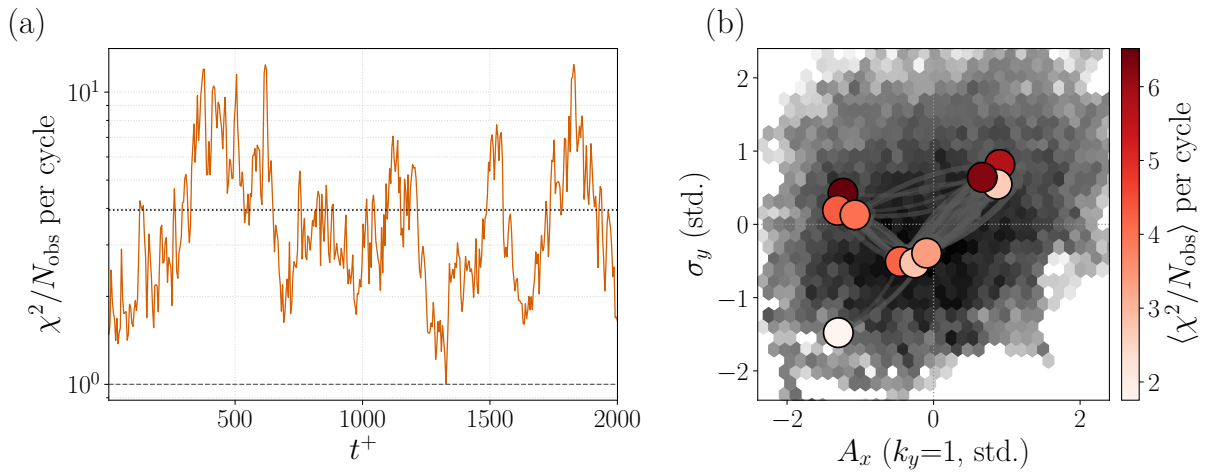


Fig. 14 Innovation consistency in time and around the cycle. (a) Per-cycle χ^2/N_{obs} over a $2000\,t^+$ run. The cycle mean (dotted) sits above consistency (dashed), the excess arriving in bursts. (b) The same statistic cycle-averaged per cluster, on the streak–vortex plane over the sampled density (grey) and coloured by $\langle\chi^2/N_{\mathrm{obs}}\rangle$. It is larger in the active, high-σ_y phase and near-consistent in the quiescent phase, tracing the cycle (faint arrows)

7 Conclusions

The near-wall self-sustaining cycle has been reduced to a differentiable latent field fitted from wall shear stress alone. Resampling the trajectory in arc length and clustering under a local Mahalanobis metric shape the basis by the attractor's geometry instead of the variance hierarchy of the decomposition. Under this geometry, the quiescent, vortex-dominated, and mature-streak phases separate into sharp populations without requiring physical labels. When integrated freely, the model stays bounded and reproduces the cycle's statistics and its invariant measure, capturing the shape of the full attractor rather than any single orbit upon it.

Since the fitted field is explicit and differentiable, it can be analysed directly as a differentiable vector field. Its slow points are located by Newton iteration and lie within the quiescent half, representing weakly unstable near-equilibria reminiscent of the cycle's lower- and upper-branch invariant solutions. Its Jacobian resolves the anatomy of the cycle phase by phase: the non-normal growth of a disturbance peaks just as the streak reaches full strength and is on the point of breaking down, whereas phase-volume contraction is strongest as the streak decays. These insights emerge directly from the fitted coefficients, without interrogating the underlying Navier–Stokes governing equations.

Assimilating wall measurements into the model quantifies the extent to which the wall determines the cycle. The wall-only model successfully reproduces the long term statistics but not the instantaneous trajectory: an intrinsic model error corresponding to an additive-inflation level of approximately one-quarter of the long-term standard deviation, which remains stable across variations in sensor noise and spatial spacing, measures the forecast error incurred from the interior motions it does not resolve. This deficit acts as a fixed fraction of the tendency at each step and concentrates, once integrated, in the active phases where the lift-up mechanism and streak breakdown are governed from off the wall. Consequently, a field fitted entirely in wall coordinates learns a conditional-mean tendency: the cycle is autonomous with respect to the wall observable in a statistical rather than a deterministic sense. These findings have been currently established for a single, moderate friction Reynolds number in a minimal channel.

Several future directions naturally follow. The interior-information deficit could be tackled using a closure model designed to inject the unresolved off-wall dynamics, which would lower this error lower bound and extend the horizon over which the wall-only forecast remains reliable. Raising the friction Reynolds number and enlarging the domain will test whether the wall remains a near-sufficient coordinate system once the near-wall cycle is no longer isolated. Incorporating wall pressure alongside shear would close a larger portion of the deficit at its source. Given that the estimator already tracks the true trajectory at 1.3× real time using sparse, noisy wall sensors, this same framework provides a direct pathway to feed near-wall control strategies.

Acknowledgements. This work used the ARCHER2 UK National Supercomputing Service (https://www.archer2.ac.uk), with resources provided by the UK Turbulence Consortium under Engineering and Physical Sciences Research Council (EPSRC) grants EP/R029326/1 and EP/X035484/1.

Funding. M.P.C. acknowledges support from the Horizon Europe Marie Skłodowska-Curie Actions Doctoral Network project SCALE (Industry empowerment to Multiphase fluid dynamics simulations using Artificial intelligence and Statistical methods on modern hardware architectures at Scale), Grant Agreement No. 101120014. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency (REA). Neither the European Union nor the granting authority can be held responsible for them. G.M.C. acknowledges the support of the EPSRC through its Doctoral Training Partnership studentship programme, grant number EP/W524608/1. A.N. acknowledges support from Schmidt Sciences, LLC.

Competing interests. The authors declare no competing interests.

Data availability. The data used to build and evaluate the model are openly available at Zenodo, https://zenodo.org/records/21629801.

Code availability. The code implementing the reduced-order model is available at https://github.com/miketwix373/RBF_LatentDynamics_MFU.

Author contributions. M. Pérez Cuadrado: Conceptualization, Methodology, Software, Formal analysis, Investigation, Writing – original draft, Visualization. G.M. Cavallazzi: Conceptualization, Methodology, Software, Writing – review & editing. A. Nóvoa: Conceptualization, Supervision, Writing – review & editing. A. Pinelli: Conceptualization, Supervision, Writing – review & editing, Funding acquisition.

References

Anderson J (2007) An adaptive covariance inflation error correction algorithm for ensemble filters. Tellus A 59(2):210–224. https://doi.org/10.1111/j.1600-0870.2006.00216.x

Andersson P, Brandt L, Bottaro A, et al (2001) On the breakdown of boundary layer streaks. J Fluid Mech 428:29–60. https://doi.org/10.1017/S0022112000002616

Aubry N, Holmes P, Lumley JL, et al (1988) The dynamics of coherent structures in the wall region of a turbulent boundary layer. J Fluid Mech 192:115–173. https://doi.org/10.1017/S0022112088001818

Bell RJ, Dean P (1970) Atomic vibrations in vitreous silica. Discuss Faraday Soc 50:55–61. https://doi.org/10.1039/DF9705000055

Benner P, Gugercin S, Willcox K (2015) A survey of projection-based model reduction methods for parametric dynamical systems. SIAM Rev 57(4):483–531. https://doi.org/10.1137/130932715

Bose ST, Park GI (2018) Wall-modeled large-eddy simulation for complex turbulent flows. Annu Rev Fluid Mech 50:535–561. https://doi.org/10.1146/annurev-fluid-122316-045241

Brandt L (2014) The lift-up effect: The linear mechanism behind transition and turbulence in shear flows. Eur J Mech B Fluids 47:80–96. https://doi.org/10.1016/j.euromechflu.2014.03.005

Brunton S, Proctor J, Kutz J (2016) Discovering governing equations from data by sparse identification of nonlinear dynamical systems. Proc Natl Acad Sci USA 113(15):3932–3937. https://doi.org/10.1073/pnas.1517384113

Butler KM, Farrell BF (1992) Three-dimensional optimal perturbations in viscous shear flow. Phys Fluids A 4(8):1637–1650. https://doi.org/10.1063/1.858386

Cavalieri AVG, Nogueira PAS (2022) Reduced-order Galerkin models of plane Couette flow. Phys Rev Fluids 7(10):L102601. https://doi.org/10.1103/PhysRevFluids.7.L102601

Cavallazzi G, Guastoni L, Vinuesa R, et al (2025) Deep reinforcement learning for the management of the wall regeneration cycle in wall-bounded turbulent flows. Flow Turbul Combust 115(3):1291–1317. https://doi.org/10.1007/s10494-024-00609-4

Cavallazzi G, Pérez-Cuadrado M, Pinelli A (2026) Offline accuracy is not enough: closed-loop instability and stabilisation of a wall-sensor neural estimator in opposition control. arXiv preprint arXiv:260630484 Under review at *Flow, Turbulence and Combustion*, arXiv:2606.30484 [physics.flu-dyn]

Chartrand R (2011) Numerical differentiation of noisy, nonsmooth data. ISRN Appl Math 2011:164564. https://doi.org/10.5402/2011/164564

Choi H, Moin P, Kim J (1994) Active turbulence control for drag reduction in wall-bounded flows. J Fluid Mech 262:75–110. https://doi.org/10.1017/S0022112094000431

Chorin AJ, Hald OH, Kupferman R (2000) Optimal prediction and the Mori–Zwanzig representation of irreversible processes. Proc Natl Acad Sci USA 97(7):2968–2973. https://doi.org/10.1073/pnas.97.7.2968

Colanera A, Magri L (2025) Quantized local reduced-order modeling in time (ql-ROM). Comput Methods Appl Mech Engrg 447:118393. https://doi.org/10.1016/j.cma.2025.118393

Colanera A, Magri L (2026) Towards extreme event prediction of turbulent flows with quantized local reduced-order models. J Phys Conf Ser 3230(1):012003. https://doi.org/10.1088/1742-6596/3230/1/012003, URL https://doi.org/10.1088/1742-6596/3230/1/012003

Colburn C, Cessna J, Bewley T (2011) State estimation in wall-bounded flow systems. part 3. the ensemble Kalman filter. J Fluid Mech 682:289–303. https://doi.org/10.1017/jfm.2011.222

Costa P (2018) A FFT-based finite-difference solver for massively-parallel direct numerical simulations of turbulent flows. Comput Math Appl 76(8):1853–1862. https://doi.org/10.1016/j.camwa.2018.07.034

Cuéllar A, Güemes A, Ianiro A, et al (2024) Three-dimensional generative adversarial networks for turbulent flow estimation from wall measurements. J Fluid Mech 991:A1. https://doi.org/10.1017/jfm.2024.432

Dee D (1995) On-line estimation of error covariance parameters for atmospheric data assimilation. Mon Weather Rev 123(4):1128–1145. https://doi.org/10.1175/1520-0493(1995)123⟨1128:OEOECP⟩2.0.CO;2

Desroziers G, Berre L, Chapnik B, et al (2005) Diagnosis of observation, background and analysis-error statistics in observation space. Q J R Meteorol Soc 131(613):3385–3396. https://doi.org/10.1256/qj.05.108

Encinar MP, Jiménez J (2019) Logarithmic-layer turbulence: A view from the wall. Phys Rev Fluids 4(11):114603. https://doi.org/10.1103/PhysRevFluids.4.114603

Evensen G (2003) The ensemble kalman filter: theoretical formulation and practical implementation. Ocean Dyn 53(4):343–367. https://doi.org/10.1007/s10236-003-0036-9

Fernex D, Noack B, Semaan R (2021) Cluster-based network modeling — from snapshots to complex dynamical systems. Sci Adv 7:eabf5006. https://doi.org/10.1126/sciadv.abf5006

Gibbs A, Su F (2002) On choosing and bounding probability metrics. Int Stat Rev 70(3):419–435. https://doi.org/10.1111/j.1751-5823.2002.tb00178.x

Gonzalez TF (1985) Clustering to minimize the maximum intercluster distance. Theor Comput Sci 38:293–306. https://doi.org/10.1016/0304-3975(85)90224-5

Guastoni L, Güemes A, Ianiro A, et al (2021) Convolutional-network models to predict wall-bounded turbulence from wall quantities. J Fluid Mech 928:A27. https://doi.org/10.1017/jfm.2021.812

Guastoni L, Rabault J, Schlatter P, et al (2023) Deep reinforcement learning for turbulent drag reduction in channel flows. Eur Phys J E 46(6):51. https://doi.org/10.1140/epje/s10189-023-00285-8

Güemes A, Discetti S, Ianiro A, et al (2021) From coarse wall measurements to turbulent velocity fields through deep learning. Phys Fluids 33(7):075121. https://doi.org/10.1063/5.0058346

Hamilton JM, Kim J, Waleffe F (1995) Regeneration mechanisms of near-wall turbulence structures. J Fluid Mech 287:317–348. https://doi.org/10.1017/S0022112095000978

Jiménez J, Moin P (1991) The minimal flow unit in near-wall turbulence. J Fluid Mech 225:213–240. https://doi.org/10.1017/S0022112091002033

Jiménez J, Pinelli A (1999) The autonomous cycle of near-wall turbulence. J Fluid Mech 389:335–359. https://doi.org/10.1017/S0022112099005066

Kaiser E, Noack B, Cordier L, et al (2014) Cluster-based reduced-order modelling of a mixing layer. J Fluid Mech 754:365–414. https://doi.org/10.1017/jfm.2014.355

Kalman R (1960) A new approach to linear filtering and prediction problems. J Basic Eng 82(1):35–45. https://doi.org/10.1115/1.3662552

Kawahara G, Kida S (2001) Periodic motion embedded in plane Couette turbulence: regeneration cycle and burst. J Fluid Mech 449:291–300. https://doi.org/10.1017/S0022112001006243

Kawahara G, Uhlmann M, van Veen L (2012) The significance of simple invariant solutions in turbulent flows. Annu Rev Fluid Mech 44:203–225. https://doi.org/10.1146/annurev-fluid-120710-101228

Kim J, Moin P, Moser R (1987) Turbulence statistics in fully developed channel flow at low Reynolds number. J Fluid Mech 177:133–166. https://doi.org/10.1017/S0022112087000892

Kravchenko AG, Choi H, Moin P (1993) On the relation of near-wall streamwise vortices to wall skin friction in turbulent boundary layers. Phys Fluids A 5(12):3307–3309. https://doi.org/10.1063/1.858692

Lloyd SP (1982) Least squares quantization in PCM. IEEE Trans Inf Theory 28(2):129–137. https://doi.org/10.1109/TIT.1982.1056489

Löfdahl L, Gad-el Hak M (1999) MEMS applications in turbulence and flow control. Prog Aerosp Sci 35(2):101–203. https://doi.org/10.1016/S0376-0421(98)00012-8

Lozano-Durán A, Bae HJ, Encinar MP (2020) Causality of energy-containing eddies in wall turbulence. J Fluid Mech 882:A2. https://doi.org/10.1017/jfm.2019.801

Lumley J (1967) The structure of inhomogeneous turbulence. In: Yaglom A, Tatarsky V (eds) Atmospheric Turbulence and Radio Wave Propagation. Nauka, Moscow, p 166–178

Moin P, Moser R (1989) Characteristic-eddy decomposition of turbulence in a channel. J Fluid Mech 200:471–509. https://doi.org/10.1017/S0022112089000741

Nagata M (1990) Three-dimensional finite-amplitude solutions in plane Couette flow: bifurcation from infinity. J Fluid Mech 217:519–527. https://doi.org/10.1017/S0022112090000829

Naguib AM, Wark CE, Juckenhöfel O (2001) Stochastic estimation and flow sources associated with surface pressure events in a turbulent boundary layer. Phys Fluids 13(9):2611–2626. https://doi.org/10.1063/1.1389284

Nakamura T, Fukami K, Hasegawa K, et al (2021) Convolutional neural network and long short-term memory based reduced order surrogate for minimal turbulent channel flow. Phys Fluids 33(2):025116. https://doi.org/10.1063/5.0039845

Naughton JW, Sheplak M (2002) Modern developments in shear-stress measurement. Prog Aerosp Sci 38(6–7):515–570. https://doi.org/10.1016/S0376-0421(02)00031-3

Nóvoa A, Magri L (2025) Online model learning with data-assimilated reservoir computers. In: Computational Science – ICCS 2025 Workshops, Lecture Notes in Computer Science, vol 15910. Springer, pp 57–64, https://doi.org/10.1007/978-3-031-97567-7_5

Nóvoa A, Racca A, Magri L (2024) Inferring unknown unknowns: Regularized bias-aware ensemble Kalman filter. Comput Methods Appl Mech Engrg 418:116502. https://doi.org/10.1016/j.cma.2023.116502

Özalp E, Nóvoa A, Magri L (2026) Real-time forecasting of chaotic dynamics from sparse data and autoencoders. Comput Methods Appl Mech Engrg 450:118600. https://doi.org/10.1016/j.cma.2025.118600

Pelleg D, Moore AW (2000) X-means: Extending k-means with efficient estimation of the number of clusters. In: Proceedings of the Seventeenth International Conference on Machine Learning (ICML). Morgan Kaufmann, pp 727–734

Pérez Cuadrado M, Cavallazzi G, Pinelli A (2026) Manifold-adapted radial basis functions for reduced-order modelling of chaotic flows. arXiv preprint arXiv:260708571

Peyré G, Cuturi M (2019) Computational optimal transport. Found Trends Mach Learn 11(5–6):355–607. https://doi.org/10.1561/2200000073

Piomelli U, Balaras E (2002) Wall-layer models for large-eddy simulations. Annu Rev Fluid Mech 34:349–374. https://doi.org/10.1146/annurev.fluid.34.082901.144919

Rempfer D (2000) On low-dimensional Galerkin models for fluid flow. Theor Comput Fluid Dyn 14:75–88. https://doi.org/10.1007/s001620050131

Robinson SK (1991) Coherent motions in the turbulent boundary layer. Annu Rev Fluid Mech 23:601–639. https://doi.org/10.1146/annurev.fl.23.010191.003125

Satopaa V, Albrecht J, Irwin D, et al (2011) Finding a "kneedle" in a haystack: Detecting knee points in system behavior. In: 2011 31st International Conference on Distributed Computing Systems Workshops (ICDCSW), IEEE, pp 166–171, https://doi.org/10.1109/ICDCSW.2011.20

Schmid PJ (2007) Nonmodal stability theory. Annu Rev Fluid Mech 39:129–162. https://doi.org/10.1146/annurev.fluid.38.050304.092139

Schwarz G (1978) Estimating the dimension of a model. Ann Stat 6(2):461–464. https://doi.org/10.1214/aos/1176344136

Sirovich L (1987) Turbulence and the dynamics of coherent structures. Part I: Coherent structures. Q Appl Math 45(3):561–571. https://doi.org/10.1090/qam/910462

Srinivasan PA, Guastoni L, Azizpour H, et al (2019) Predictions of turbulent shear flows using deep neural networks. Phys Rev Fluids 4(5):054603. https://doi.org/10.1103/PhysRevFluids.4.054603

Tandeo P, Ailliot P, Bocquet M, et al (2020) A review of innovation-based methods to jointly estimate model and observation error covariance matrices in ensemble data assimilation. Mon Weather Rev 148(10):3973–3994. https://doi.org/10.1175/MWR-D-19-0240.1

Toh S, Itano T (2003) A periodic-like solution in channel flow. J Fluid Mech 481:67–76. https://doi.org/10.1017/S0022112003003768

Trefethen LN, Trefethen AE, Reddy SC, et al (1993) Hydrodynamic stability without eigenvalues. Science 261(5121):578–584. https://doi.org/10.1126/science.261.5121.578

Vlachas PR, Pathak J, Hunt BR, et al (2020) Backpropagation algorithms and reservoir computing in recurrent neural networks for the forecasting of complex spatiotemporal dynamics. Neural Netw 126:191–217. https://doi.org/10.1016/j.neunet.2020.02.016

Waleffe F (1997) On a self-sustaining process in shear flows. Phys Fluids 9(4):883–900. https://doi.org/10.1063/1.869185

Waleffe F (2001) Exact coherent structures in channel flow. J Fluid Mech 435:93–102. https://doi.org/10.1017/S0022112001004189

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- ga.pdf