

Searching turbulence for periodic orbits with dynamic mode decomposition

Jacob Page^{1,†} and Rich R. Kerswell¹

¹DAMTP, Centre for Mathematical Sciences, University of Cambridge, Cambridge CB3 0WA, UK

(Received 4 June 2019; revised 11 December 2019; accepted 20 December 2019)

We present a new method for generating robust guesses for unstable periodic orbits (UPOs) by post-processing turbulent data using dynamic mode decomposition (DMD). The approach relies on the identification of near-neutral, repeated harmonics in the DMD eigenvalue spectrum from which both an estimate for the period of a nearby UPO and a guess for the velocity field can be constructed. In this way, the signature of a UPO can be identified in a short time series without the need for a near recurrence to occur, which is a considerable drawback to recurrent flow analysis, the current state of the art. We first demonstrate the method by applying it to a known (simple) UPO and find that the period can be reliably extracted even for time windows of length one quarter of the full period. We then turn to a long turbulent trajectory, sliding an observation window through the time series and performing many DMD computations. Our approach yields many more converged periodic orbits (including multiple new solutions) than a standard recurrent flow analysis of the same data. Furthermore, it also yields converged UPOs at points where the recurrent flow analysis flagged a near recurrence but the Newton solver did not converge, suggesting that the new approach can be used alongside the old to generate improved initial guesses. Finally, we discuss some heuristics on what constitutes a ‘good’ time window for the DMD to identify a UPO.

Key words: low-dimensional models, computational methods

1. Introduction

Since the discovery of the first unstable periodic orbit (UPO) in a transiently turbulent flow by Kawahara & Kida (2001), there has been a surge in interest and a large number of other periodic solutions found both in this configuration (e.g. Viswanath 2007; Cvitanovic & Gibson 2010) and in other canonical turbulent flows (Chandler & Kerswell 2013; Willis, Cvitanovic & Avila 2013; Lucas, Caulfield & Kerswell 2017). The discovery of a large number of UPOs supports a perspective of turbulence in which the flow is viewed as a trajectory in a very high-dimensional dynamical system, wandering between unstable exact coherent structures (Kerswell 2005; Kawahara, Uhlmann & van Veen 2012). Individual UPOs can offer a good deal of insight into the physical processes sustaining the turbulent flow owing to their simple time dependence (Waleffe 1997; Wang, Gibson & Waleffe 2007; Hall &

† Email address for correspondence: jacob.page@damtp.cam.ac.uk

Sherwin 2010). Furthermore, periodic orbit theory (Cvitanović *et al.* 2016) suggests that statistics of the turbulence can be predicted from the UPOs if enough of them are found (see the attempts in Chandler & Kerswell (2013) and Lucas & Kerswell (2015)). Despite this increasing interest, the methods for finding UPOs are somewhat crude and have not changed significantly in the two decades since the first turbulent solutions were discovered.

The standard method for finding and converging UPOs begins with a search for near recurrences in data from a numerical simulation. In such a ‘recurrent flow analysis’, likely UPOs are flagged when the distance in state space (measured with an L_2 norm) between the present and past states drops below a threshold value (Kawahara & Kida 2001; Viswanath 2007; Cvitanovic & Gibson 2010; Chandler & Kerswell 2013). The resulting set of candidate orbits, each augmented with a guessed period from the time between two similar states, is then input into a Newton solver. The main downside of the approach is that it requires the turbulence to shadow a periodic orbit for at least one full cycle, and hence it can be increasingly ineffective as the Reynolds number is increased (Chandler & Kerswell 2013). In addition, the sensitivity of the Newton solver to initial conditions can result in failures at the convergence stage when there may in fact be a UPO nearby. In this paper we introduce a new method based on dynamic mode decomposition (DMD) that goes some way to addressing these issues.

The DMD algorithm was originally invented by Schmid (2010) as a post-processing tool for numerical and experimental data. It can also serve as an excellent data-driven alternative for demanding computational tasks, for example global stability calculations in complex geometries (if the data were generated from the linearised equations of motion, see Schmid *et al.* (2011)), since it requires only raw data in the form of snapshot pairs. The output of the algorithm is a linear operator that (in a least squares sense over the input data) provides the best mapping between snapshots separated by a fixed time δt . An eigendecomposition of this operator yields a set of ‘dynamic modes’ which evolve exponentially in time, the time dependence being set by the associated eigenvalues. Crucially, this approach can identify frequencies of oscillation with corresponding periods that can be far longer than the time window over which the observations were recorded (Rowley *et al.* 2009; Schmid 2010; Schmid *et al.* 2011). Since its invention, a number of variants of the algorithm have been proposed and DMD has found applications in areas of science and engineering beyond fluid mechanics (Jovanović, Schmid & Nichols 2014; Williams, Kevrekidis & Rowley 2015; Kutz *et al.* 2016).

Despite its apparent linearity, DMD can also be a useful tool in strongly nonlinear flows. Our interest in applying DMD to turbulent data as a means to identify UPOs rests on the equivalence (under strict conditions) of the algorithm to the Koopman decomposition (Rowley *et al.* 2009; Tu *et al.* 2014; Williams *et al.* 2015) of a nonlinear dynamical system. The Koopman operator is an infinite-dimensional, linear operator that propagates observables of the state forward in time along trajectories of the dynamical system (Koopman 1931; Mezić 2005), and a Koopman decomposition of the state leads to a representation in which the nonlinear dynamics is expressed as a linear superposition of Koopman modes (e.g. fixed velocity fields) which evolve exponentially in time. Koopman decompositions have been performed analytically for some simple nonlinear ordinary (Bagheri 2013; Brunton *et al.* 2016; Rowley & Dawson 2017) and partial differential equations (Page & Kerswell 2018).

Applying a Koopman decomposition to a turbulent flow yields a representation of the state as a superposition of a set of harmonic averages and a broadband continuous spectrum (Mezić & Banaszuk 2004; Mezić 2005, 2013; Arbabi & Mezić 2017),

but individual simple invariant sets also possess their own local Koopman decompositions. For example, Mezic (2019) has shown that the Koopman eigenvalues for a nonlinear system collapsing onto a limit cycle consists of a set of repeated neutral harmonics (the limit cycle's fundamental frequency and higher harmonics) and an infinite lattice of decaying eigenvalues which can be determined from linear combinations of the cycle's Floquet multipliers. However, in systems with multiple exact coherent structures there are multiple local Koopman decompositions around each structure. These linear representations break down (are no longer convergent) at certain cross-over points in state space (Page & Kerswell 2019), a fact which impacts the ability of DMD to extract the Koopman modes associated with the expansions. That DMD no longer coincides with Koopman theory when the observation window includes the cross-over point is likely the reason behind the poor performance of the algorithm when Bagheri (2013) applied it to the 'spin up' problem of flow past a cylinder. Page & Kerswell (2019) demonstrated the phenomenon explicitly in a Stuart–Landau equation and numerically along heteroclinic connections in the Navier–Stokes equations. In doing this, it was observed that DMD can extract some of the Koopman eigenvalues associated with the local Koopman decomposition around a periodic orbit – although only if the computation is performed within the 'Koopman expansion zone' where the local decomposition is convergent. Here, we go further and layout how DMD can be applied to turbulent data to identify and converge UPOs.

The remainder of this paper is organised as follows. In § 2 we outline a procedure for performing and post-processing DMD calculations to generate guesses for periodic orbits and demonstrate its utility by applying it to a gently periodic edge state. In § 3 the method is then used to search for periodic orbits close to a long turbulent trajectory and is compared to recurrent flow analysis. We examine some of the DMD guesses further in § 4, and identify characteristics of 'good' initial guesses. Finally, concluding remarks are provided in § 5.

2. Methodology

In this paper we present a new method for identifying periodic orbits using DMD applied to short turbulent trajectories and contrast it to recurrent flow analysis, the current state of the art. To do this we post-process numerical simulation data using both methods and attempt to converge the resulting UPO guesses via Newton–Krylov iteration. Throughout, the flow configuration is plane Couette flow in a minimal flow unit (Hamilton, Kim & Waleffe 1995) non-dimensionalised by the upper/lower plate speed U_0 and the channel half-height d . The Reynolds number is fixed at $Re := U_0 d / \nu = 400$, while the exact horizontal dimensions of the periodic computational domain are $(L_x, L_y) = (1.755\pi, 1.2\pi)$ to match Kawahara & Kida (2001).

In our direct numerical simulations, the incompressible Navier–Stokes equations are solved using a fractional-step method with an implicit Crank–Nicholson scheme for the diffusive terms and a third-order Runge–Kutta scheme for the advection terms (Dubief *et al.* 2005). Second-order finite differences are used in all three spatial directions with grid stretching applied in the vertical. For the majority of results the grid resolution is $N_x \times N_y \times N_z = 96 \times 96 \times 129$. Note that the slight difference in box size from the $L_x = 1.75\pi$ box used by Cvitanovic & Gibson (2010) (and the solutions catalogued at channelflow.org) precludes a direct comparison with this extensive set of UPOs. Altering our box size to 1.75π indicates that some of UPOs

we find are known solutions listed at channelflow.org. At our resolution we can match the periods of these known periodic orbits to within a relative error of $\sim 0.1\%$, and have verified that this error can be reduced by further increasing our resolution. We make every effort to connect our results to previously discovered exact solutions (see table 1).

In a recurrent flow analysis (e.g. Chandler & Kerswell 2013), guesses for periodic orbits are identified via local minima in an L_2 norm,

$$R(t, T) := \frac{\|\mathbf{u}(t+T) - \mathbf{u}(t)\|}{\|\mathbf{u}(t)\|}, \quad (2.1)$$

which fall below some threshold R_{thresh} . If the minima occur at times t_i , then the guesses for periodic orbits are simply the corresponding states $\{\mathbf{u}(t_i)\}$, each augmented with a guessed period $T_g = T_i$ – the future time at which the local minimum in $R(t, T)$ occurs. Note that in the form (2.1) we are comparing future states at $t+T$ with present states at t in order to allow for comparison to results from DMD with observation windows starting at time t .

As discussed in the introduction, there are shortcomings to this approach for generating guesses for periodic orbits. For instance, given a near recurrence, the (future) time T at which the minimum value of $R(t, T)$ occurs may differ significantly from the period of the nearby UPO. It is plausible that the sensitivity of Newton–Raphson to initial conditions could result in a lack of convergence even if there is indeed a structure nearby. In these situations, DMD may be a useful tool since it can be used to extract a set of modes associated with ‘background’ periodic motion, resulting in a more robust initial condition for the root finder and a more informed estimate of the period. Perhaps more importantly, DMD does not require a near recurrence to occur before it can ‘sense’ periodic motion since the algorithm allows for the extraction of frequencies corresponding to oscillations far longer than the length of observation window (Schmid 2010). Both of these scenarios are sketched in figure 1.

In DMD, vector observables of the state at discrete times $\{t_j\}$ are stored in a data matrix (Schmid 2010; Tu *et al.* 2014),

$$\Psi^t = [\psi(\mathbf{u}(t_1)) \quad \psi(\mathbf{u}(t_2)) \quad \cdots \quad \psi(\mathbf{u}(t_M))], \quad (2.2)$$

where the observable functionals that make up the elements of ψ are a design choice, often motivated by knowledge of the underlying dynamical system (e.g. Williams *et al.* 2015). A second data matrix is formed with observations made a time δt (another design choice) later,

$$\Psi^{t+\delta t} = [\psi(\mathbf{u}(t_1 + \delta t)) \quad \psi(\mathbf{u}(t_2 + \delta t)) \quad \cdots \quad \psi(\mathbf{u}(t_M + \delta t))], \quad (2.3)$$

where we refer to corresponding columns in Ψ^t and $\Psi^{t+\delta t}$ as a ‘snapshot pair’. Note that the observation times $\{t_j\}$ need not be uniformly spaced in time. Furthermore, the ordering of the columns in Ψ^t can be random and may span multiple datasets; all that is required is that snapshot pairs are observations from the same trajectory and that the separation time δt is fixed across all snapshot pairs (Tu *et al.* 2014).

The DMD operator is the best-fit (in a least-squares sense) linear operator that maps between the data matrices,

$$\hat{K} := \Psi^{t+\delta t} (\Psi^t)^+, \quad (2.4)$$




FIGURE 1. Cartoon of a periodic orbit (black) and two trajectories where DMD provides a useful alternative to recurrent flow analysis. The orange line is a segment of an orbit which does not shadow the UPO for a full period and hence is useless for recurrent flow analysis. The blue line does shadow the UPO for a full cycle, but using the local minimum in $R(t, T)$ may yield an inaccurate estimate of the period. The shaded region identifies the ‘Koopman expansion zone’ around the UPO, discussed in the text.

where the superscript $+$ identifies the Moore–Penrose pseudo-inverse, accomplished via a singular value decomposition (SVD) of the data matrix $\Psi^t = U \Sigma W^H$. In reality, the dimensionality of the problem requires that \hat{K} is never computed directly, but instead the $r \times r$ (where r is the rank of Ψ^t) matrix $U^H \hat{K} U$ is used – the projection of the DMD matrix onto POD (proper orthogonal decomposition) modes.

To search for periodic orbits we will be applying DMD to relatively short turbulent trajectories ($O(100)$ advective time units with $M \sim 100$ snapshot pairs), and hence the data matrix will likely be full rank $r = M$. This tends to lead to overfitting and unphysical DMD modes featuring sharp gradients and large-amplitude, localised features as the DMD algorithm attempts to fit exponential dynamics to all velocity scales. These aphysical DMD modes can be avoided through a variety of techniques, e.g. by either performing a low-rank truncation or by introducing a sparsity constraint (Jovanović *et al.* 2014). However, rather than selecting *a priori* a threshold singular value/energy for a low-rank projection, or modifying the DMD algorithm, we instead attempt to find a DMD operator that can map between velocity fluctuations above a certain threshold by adding a small amount of uniform white noise to our observable vector,

$$\psi(\mathbf{u}) := \mathbf{u} - \mathbf{u}_C + \delta \mathbf{w}_\epsilon, \tag{2.5}$$

where $\mathbf{u}_C := z \mathbf{e}_x$ is the laminar base state and $\delta \mathbf{w}_\epsilon$ is uniform white noise, where each component is uniformly distributed $(\delta \mathbf{w}_\epsilon)_j \stackrel{\text{iid}}{\sim} U[-\epsilon, \epsilon]$. The introduction of white noise results in a clear cutoff in the singular values which we use to perform a low-rank projection, with $\sigma_j \sim \sigma_{\min} \approx \text{const.}$ beyond a threshold mode number j (see figure 2). The minimum singular value σ_{\min} scales proportionally with ϵ . Throughout, we set $\epsilon = 10^{-7}$, although we have verified that our results (DMD eigenvalues and modes) are robust to modest changes in ϵ of at least plus or minus a decade.

Once the (projected) DMD operator is computed, an eigendecomposition can be performed. Under certain conditions (see Williams *et al.* 2015; Rowley & Dawson




FIGURE 2. Singular values of the data matrix Ψ^t for various white noise amplitudes ϵ in an example computation with $M = 100$ snapshots. (Red) $\epsilon = 10^{-5}$, (green) $\epsilon = 10^{-7}$, (blue) $\epsilon = 10^{-9}$, (cyan) $\epsilon = 0$. Note the flattening of the singular values at some $\sigma_{min}(\epsilon)$.

2017; Page & Kerswell 2019) the output of DMD coincides with a Koopman analysis. In these cases, approximate Koopman modes for the observable $\psi(\mathbf{u})$ can be obtained as right eigenvectors of $\hat{\mathbf{K}}$ (DMD modes),

$$\hat{\mathbf{K}}\mathbf{v}_j = e^{\lambda_j \delta t} \mathbf{v}_j. \tag{2.6}$$

Note that if \mathbf{v}' is a right eigenvector of $\mathbf{U}^H \hat{\mathbf{K}} \mathbf{U}$, the DMD mode is obtained from $\mathbf{v} = \mathbf{U}\mathbf{v}'$ (Schmid 2010). Koopman eigenfunctions are related to the left eigenvectors of $\hat{\mathbf{K}}$, $\{\mathbf{w}_j\}$, via

$$\varphi_{\lambda_j}(\mathbf{u}) = \mathbf{w}_j^H \psi(\mathbf{u}), \tag{2.7}$$

which can be connected to the left eigenvector \mathbf{w}'_j of $\mathbf{U}^H \hat{\mathbf{K}} \mathbf{U}$ via $\mathbf{w}_j = \mathbf{U}\mathbf{w}'_j$. Koopman eigenfunctions are special scalar observables of the state, \mathbf{u} , that have an exponential dependence on time (Mezić 2005). To see this in (2.7), consider the same equation evaluated a time δt later along a trajectory,

$$\begin{aligned} \mathbf{w}_j^H \psi(\mathbf{u}(t + \delta t)) &\approx \mathbf{w}_j^H \hat{\mathbf{K}} \psi(\mathbf{u}(t)), \\ &= e^{\lambda_j \delta t} \mathbf{w}_j^H \psi(\mathbf{u}(t)). \end{aligned} \tag{2.8}$$

For more on the connection between DMD and Koopman, and the requirements that must be satisfied to accurately extract Koopman eigenfunctions, see Williams *et al.* (2015), Rowley & Dawson (2017) and Page & Kerswell (2019).

2.1. Spotting a periodic orbit

In a dynamical system with a single attracting periodic orbit, Mezić (2019) has shown that the spectrum of the associated Koopman operator consists of a set of repeated harmonics (multiples of the fundamental frequency of the limit cycle) along with

an infinite lattice of decaying eigenvalues that can be computed from the Floquet multipliers. In the turbulent Couette flow considered in this work, all periodic orbits are unstable and each orbit will have an associated locally convergent Koopman decomposition (Page & Kerswell 2019) of a similar form to that derived by Mezic (2019) but with unstable eigenvalues related to the set of unstable Floquet multipliers. While DMD on a trajectory shadowing the UPO for (potentially less than) one cycle is unlikely to be able to identify the stability properties of the nearby structure, the results presented in Page & Kerswell (2019) indicate that the signature of the periodic orbit – its neutral harmonics – in the DMD results can remain even for relatively short time windows.

In previous studies involving DMD on flows collapsing onto limit cycles (e.g. Bagheri 2013; Page & Kerswell 2019), identifying the underlying period in the DMD eigenvalue spectrum is straightforward either because the period is known beforehand or the DMD can be run over an indefinitely long time horizon due to the stability of the exact coherent structure. However, in a turbulent flow it is unknown *a priori* at any given time whether we are shadowing a periodic orbit, and if so what its period is. Therefore, given an eigenvalue spectrum from a DMD calculation on a segment of a turbulent orbit we would like to both (i) estimate the period of a nearby UPO (if one exists) and (ii) assess the quality of our ‘guess’ to make an informed decision about whether there is indeed a nearby structure.

For step (i), the estimation of a UPO period, we proceed as follows: DMD of a truly periodic solution should produce a set of eigenvalues which all have $\lambda_r = 0$, hence a maximum cutoff growth/decay rate, λ_r^{max} , is chosen to select a subset of the DMD modes, $\tilde{\Lambda} := \{\lambda \mid |\text{Re}(\lambda)| \leq \lambda_r^{max}\}$. The truncated set of eigenvalues $\tilde{\Lambda}$ contains at least one mode with $\text{Im}(\lambda) = 0$ and pairs of complex conjugate eigenvalues. Of the modes with $\text{Im}(\lambda) = 0$, the one with the smallest real component is assumed to be the time average of a nearby UPO, while the first n modes with positive imaginary components are assumed to constitute the UPO’s fundamental frequency and its $n - 1$ higher harmonics. Under this assumption, a fundamental frequency can be estimated via

$$\omega_f(n) := \frac{2}{n(n + 1)} \sum_{j=1}^n \omega_j, \tag{2.9}$$

where $\omega_j := \text{Im}(\lambda_j)$. The guess for a velocity field at a point on the UPO is then built from the time average mode and the first n complex conjugate pairs (this procedure is described in more detail in § 2.3 below), while the period is estimated as $T_g = 2\pi/\omega_f$. For step (ii), the estimate of how good such a guess is expected to be, we determine the degree to which the finite-frequency modes approximate a series of harmonics,

$$\varepsilon_\omega(n) := \frac{1}{n|\omega_f|^2} \sum_{j=1}^n |\omega_j - j\omega_f|^2, \tag{2.10}$$

with $\varepsilon_\omega = 0$ for a perfectly periodic signal. In the following analysis, DMD results for which values of ε_ω fall below a prescribed threshold are taken to indicate the likely existence of a nearby UPO.

2.2. Motivating example: gently periodic edge state

To demonstrate the mechanics of the approach outlined above we first apply it to a trajectory that shadows the simplest UPO in this box (the edge state) for several

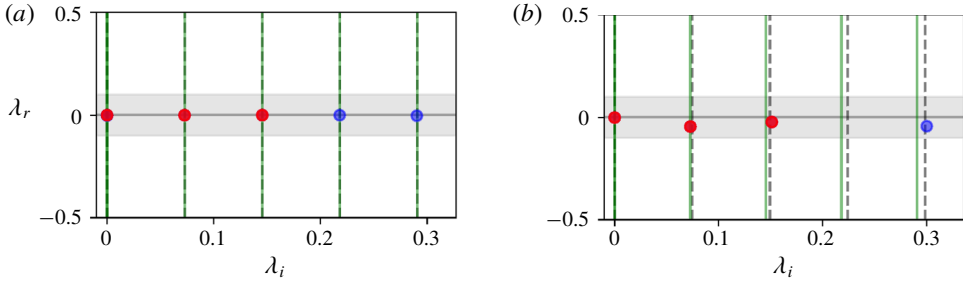


FIGURE 3. DMD eigenvalues obtained from observations on the ‘edge trajectory’ described in the text. (a) Observation window length matches the period of the edge UPO, $T_w = T_\omega$. (b) Observation window is a quarter of the UPO period, $T_w = T_\omega/4$. Green lines identify integer multiples of the fundamental frequency of the nearby UPO, $2\pi/T_\omega$. Dashed black lines correspond to the fundamental frequency estimated from the first two harmonics (highlighted red) in the DMD via (2.9). The shaded region indicates a threshold $\lambda_r^{\max} = 0.1$. In both calculations $\delta t = 1$, $M = 50$.

complete cycles before becoming turbulent. This ‘gentle’ periodic orbit was first discovered by Kawahara & Kida (2001) and has a period $T_\omega \approx 86.4$ in this box (see also Kawahara (2005)). The trajectory is obtained via repeated bisection between relaminarising and turbulent trajectories (Schneider, Eckhardt & Yorke 2007) and we store one full period at a resolution of $\Delta t = 0.1$ (five times coarser than the simulation timestep). This choice yields ~ 865 possible observation times $\{t_j\}$ for the data matrix Ψ^t ; the exact number, $865 - \delta t/\Delta t$, depends on the snapshot pair separation time δt . In our DMD calculations we select $\delta t = 10\Delta t = 1$ as the separation time between corresponding snapshots in Ψ^t and $\Psi^{t+\delta t}$ and randomly draw $M = 50$ snapshot pairs from a possible 855.

DMD spectra from two calculations on this dataset are reported in figure 3. The two calculations are identical in all respects apart from the length of the observation window, T_w , over which the DMD is performed. In the first case, the length of the observation window matches the period of the underlying UPO, $T_w \approx T_\omega$, and the DMD spectrum consists of a single neutral mode and a series of near-perfect harmonics. The fundamental frequency estimated from the first purely imaginary mode and the first harmonic matches the fundamental frequency of the UPO, while the associated value ε_ω is very close to zero. This result is entirely expected owing to the observation of a full period and the proximity of the trajectory to the UPO. Note that past studies of Koopman analysis around periodic orbits (e.g. Mezic 2019; Page & Kerswell 2019) indicate we should also expect to find unstable eigenvalues in the DMD results, $\lambda_{nm} = n\sigma + im\omega_f$ with $n, m \in \mathbb{N}$, where $\exp(\sigma T_\omega)$ is the single unstable Floquet multiplier of the edge state. However, approximations to these eigenvalues are only obtained over much longer observation windows, which run for several periods as the trajectory moves away from the edge (not shown here, but see the limit cycle in Page & Kerswell (2019)).

Interestingly, the results of the DMD calculation on the window $T_w = T_\omega/4$ also clearly show the signature of the UPO despite the short observation time. The eigenvalues are now slightly off the neutral line $\lambda_r = 0$, but there are still clearly repeated harmonics (the third harmonic $\sim 3\omega_f$ is missing). The period estimated from the first two harmonics is very close to the true period of the UPO, and the corresponding value of $\varepsilon_\omega \lesssim 10^{-3}$.




FIGURE 4. Streamwise velocity on a plane midway between the two walls for the (real part of the) DMD modes corresponding to the eigenvalue spectra reported in figure 3. The neutral mode and the first and second harmonics are shown. (a) $T_w = T_\omega$, (b) $T_w = T_\omega/4$. Note modes have been normalised so that contours run -1 (black) to $+1$ (bright yellow).

The two DMD calculations are compared further in figure 4. The neutral DMD mode obtained over the short time window $T_\omega/4$ is qualitatively similar to the true neutral Koopman mode associated with the orbit (the DMD mode from the observation window $T_w = T_\omega$), although the performance is notably worse for the higher harmonics. However, initial conditions built from these modes (via the procedure described in the following section) still rapidly converge to the UPO when used as an initial guess in Newton–Krylov routine.

The performance of shortened observation windows on the edge UPO is examined for a wide range of window lengths, each averaged over a large range of possible start times $t_s \in [0, T_\omega)$, in figure 5. The results indicate that short windows $T_w \lesssim 0.4T_\omega$ can perform well sometimes. On average, the DMD gets progressively worse at identifying the underlying UPO as the length of the observation window is reduced, although the accuracy of an individual computation depends subtly on the start time of observations t_s as this controls what segment of the UPO is seen by the DMD. Importantly, the good results shown for the $T_w = T_\omega/4$ window in figure 3 would not have been obtained if an alternate start time had been selected. General rules of thumb as to what exactly the DMD needs to see in order to provide a robust estimate of the period T_ω of a UPO are discussed in §4 following the presentation of results from a turbulent computation.




FIGURE 5. (a) Average predicted period $T_g = 2\pi/\omega_f$ of the edge-state UPO as a function of observation window length using estimates based on two (red) and three (blue) harmonics. For each value of T_w , averaging is performed over sets of continuous time windows with start times separated by $\Delta t_s = 5$. Shaded regions identify ± 1 sample standard deviation. The horizontal dashed line identifies the true period $T_\omega = 86.4$. (b) ‘Quality’ of the guessed period as determined by (2.9).

2.3. Building a guess for a periodic orbit

While the above results were obtained by examining trajectories that shadow a known orbit, the periodic orbits close to a turbulent trajectory will not be known beforehand, and we would like to also use the output of DMD to build initial ‘guesses’ to supply to a Newton solver if the value of ε_ω drops below the prescribed threshold. The approach we adopt is outlined here.

For DMD results where $\varepsilon_\omega < \varepsilon_\omega^{thresh}$ for a given number of harmonics n , the period of the UPO is estimated from the fundamental frequency $\omega_f(n)$, defined in (2.9), by $T_g = 2\pi/\omega_f$. The neutral mode and the first n complex conjugate pairs used to estimate $\omega_f(n)$ are then used to construct a guess for the velocity field on the UPO,

$$\mathbf{u}_g(t) = z\mathbf{e}_x + \sum_{j=-n}^n a_j \mathbf{v}_j e^{\lambda_j t}. \tag{2.11}$$

There are a variety of ways to define the unknown coefficients $\{a_j\}$. We set them such that the average (squared) deviation of the guess from the turbulent trajectory used in the DMD is smallest, i.e. by minimising

$$J(\mathbf{a}) := \frac{1}{M} \sum_{m=0}^{M-1} \left| \boldsymbol{\psi}(\mathbf{u}(x, m\delta t)) - \sum_{j=-n}^n a_j \mathbf{v}_j e^{m\lambda_j \delta t} \right|^2. \tag{2.12}$$

The minimising $\hat{\mathbf{a}}$ has been presented previously in Page & Kerswell (2019). Finally, the time t in (2.11) is selected as $\hat{t} = \arg \min_t |\mathbf{u}_g(t) - \mathbf{u}(t)|$.

3. Turbulent orbits

Motivated by the success of the DMD approach at identifying properties of the edge-state UPO on short time windows, we now apply our method to turbulent data in the same flow configuration and attempt to converge UPO guesses generated as

described above with a Newton solver. To allow for a direct comparison to recurrent flow analysis, we perform many DMD computations on a short time window which is passed through a continuous turbulent trajectory of length >1000 advective time units.

3.1. Recurrent flow analysis

The results of a recurrent flow analysis are summarised in figure 6(a), where we report contours of $R(t_s, T)$ (see (2.1)). Local minima which fall below a threshold value of $R_{\text{thresh}} = 0.07$ are supplied as guesses in a Newton–Krylov algorithm. Of the 13 guesses highlighted, 6 converged (green +, while red \times indicate failures).

It is notable that the guessed period, determined from the T at which the local minimum in $R(t_s, T)$ occurs, can often differ significantly from the period of the converged solution (note the vertical distance between green + and green \circ). For example, the (good) guesses flagged at $t = 235$ and $t = 262$ converged to periodic orbits with relative errors in the period against the guess of 12.3% and 10.5% respectively. As we shall see below, the combination of the state $\mathbf{u}(t)$ and period T which correspond to a local minimum in $R(t, T)$ can often result in failure in the Newton algorithm even if there is a nearby UPO. This is particularly evident in figure 6 for $t > 800$, where recurrent flow analysis fails to identify and successful guesses, while our new approach results in three converged UPOs.

3.2. Dynamic mode decomposition

The output of the DMD-based UPO detection algorithm outlined in § 2.1 is shown directly below the recurrence plot in figure 6, with guessed periods T_g reported based on both two and three harmonics. To generate these results, DMD computations were performed with a fixed observation window of length $T_w = 60$, which was moved through the time series in steps of $\Delta_w = 5$. Highlighted points along the curves $T_g(t_s)$ indicate that the value of ε_ω (2.10) dropped below a specified threshold. These points tend to be arranged in clusters with similar guesses for the period T_g , consistent with the trajectory briefly shadowing a UPO before being flung out of a particular Koopman expansion zone (Page & Kerswell 2019). The detection threshold was set at $\varepsilon_\omega = 10^{-3}$ for $n = 2$ and $\varepsilon_\omega = 5 \times 10^{-3}$ for $n = 3$; the threshold was increased for $n = 3$ since resolution of the third harmonic is more challenging. The increase in the number of harmonics used in the search for UPOs results in a drop in the number of proposed guesses, most notably for guesses with long periods. For example, the number of guesses with $T_g > 80$ falls from 16 to 3. However, in general there is good correspondence for the locations $\{t'_s\}$ of the possible UPOs between the $n = 2$ and $n = 3$ results.

The choice of observation window length $T_w = 60$ was motivated by two factors: (i) if the time window is too long and the turbulence visits the neighbourhoods of multiple UPOs, the spectrum from the DMD will likely be unrelated to any of them (see the discussion in the introduction and Page & Kerswell (2019)) and (ii) $T_w = 60$ is shorter than many of the known UPOs in this box, most of which have periods $T_\omega > 75$ (Cvitanovic & Gibson 2010) and will test our method's ability to identify and converge solutions without a near recurrence. One immediate consequence of the choice $T_w = 60$ is a loss of resolution of the longest UPO found by the recurrent flow analysis at $t_s = 412$ with period $T_\omega = 103.01$, although notably the DMD-based algorithm does flag a pair of likely UPOs with different periods close to this point – a behaviour which is discussed in more detail below. However, in general DMD

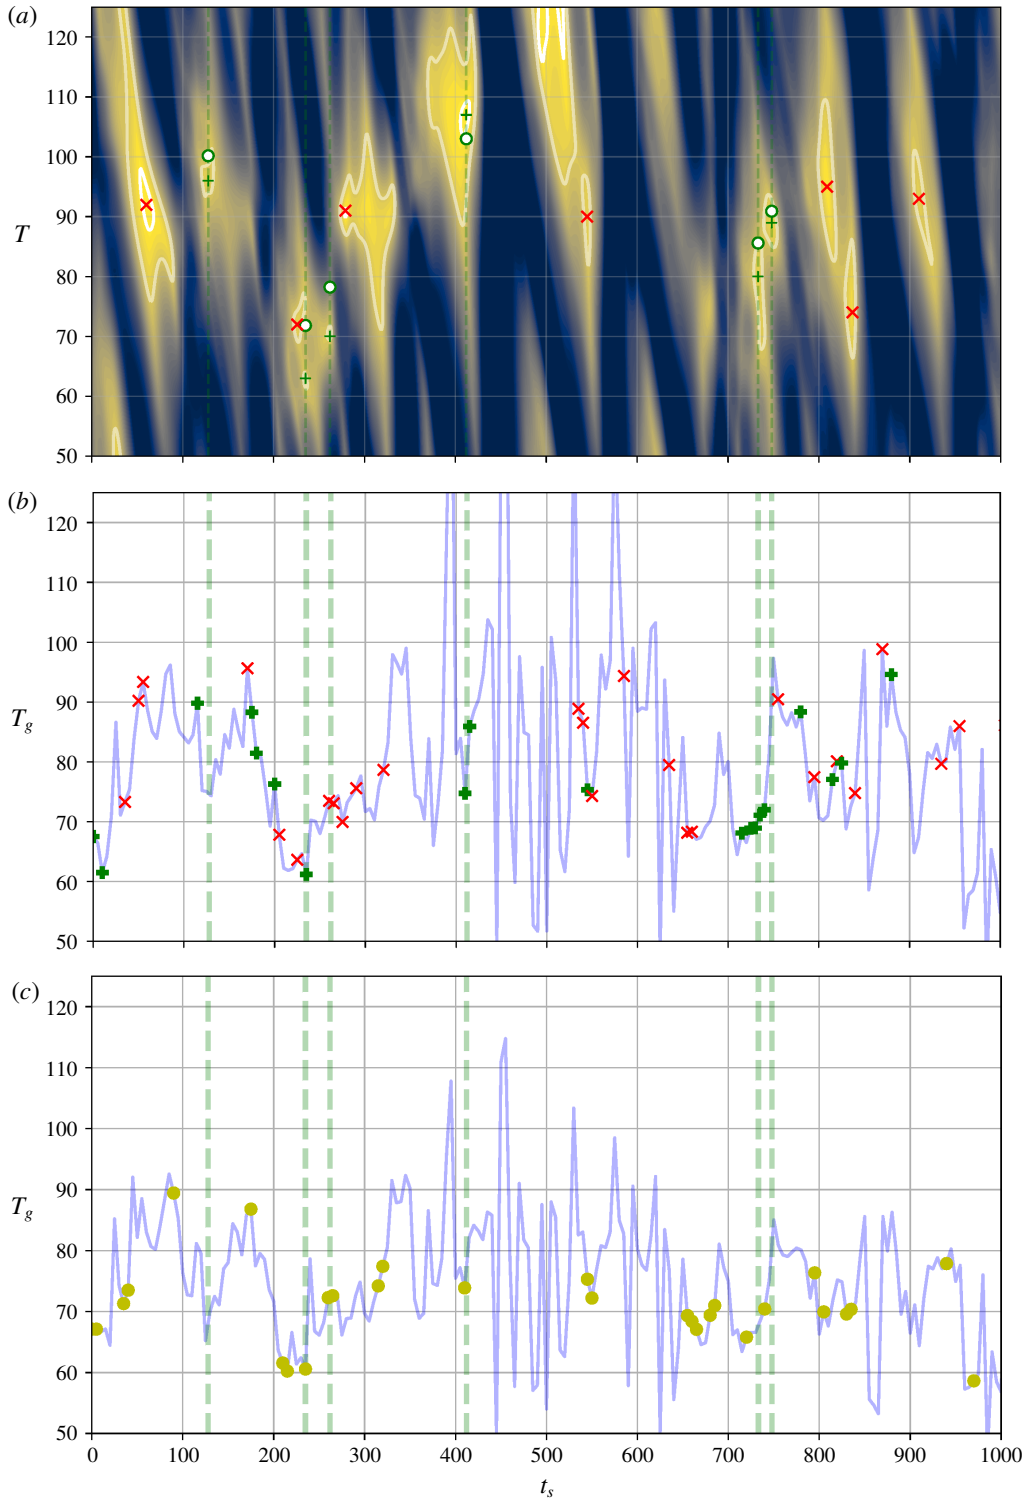


FIGURE 6. For caption see next page.

FIGURE 6 (cntd). Summary of UPO guesses generated via recurrent flow analysis (a) and the DMD procedure outlined in the text performed on time windows of length $T_w = 60$ using either two (b) or three (c) harmonics to estimate a period T_g (blue line). In the recurrence plot, green crosses identify successful guesses, while filled circles indicate the period of the converged solution. Red crosses are guesses that failed to converge. Note that the solid white contours identify values of $R = 0.05$ and $R = 0.07$. DMD results are plotted against the start time of the observations, t_s , with a time difference $\Delta_w = 5$ separating subsequent calculations. $M = 100$ snapshot pairs were used in each calculation, with $\delta t = 1$. Points highlighted by symbols on the curves indicate where the value ε_ω dropped below a prescribed threshold ($\varepsilon_\omega = 10^{-3}$ for two harmonics and $\varepsilon_\omega = 5 \times 10^{-3}$ for three harmonics). All points satisfying this threshold in the two harmonic case were then fed into a Newton–Raphson algorithm: bold green crosses identify convergences, red crosses are failures. Yellow spots in the lowest plot indicate where the three-harmonic procedure flagged likely UPOs (no attempt was made to converge these).

identifies guesses with similar periods to the UPOs spotted in the recurrent flow analysis. Some heuristics around characteristics possessed by ‘good’ DMD guesses are discussed in §4, although ultimately we suspect the selection of T_w in new applications is something that will have to be determined via trial-and-error subject to the tradeoff described above.

All of the guesses in figure 6 identified by DMD in the two-harmonic case were supplied as initial conditions in a Newton–Krylov routine. Of the 46 guesses highlighted, 20 converged to periodic orbits. This total includes UPOs found in regions where the recurrent flow analysis indicated the presence of a structure, but also in regions where no near-recurrent episodes were detected; note the guesses for $t_s \lesssim 50$ and for $150 \lesssim t_s \lesssim 200$. The various UPOs converged from the DMD-based guesses identified in figure 6 are summarised in table 1, and we will now discuss differences between these results and the recurrent flow analysis in some detail. Our discussion will explore (i) cases where a UPO is found using recurrent flow analysis but the DMD-generated guess fails to converge; (ii) cases where the DMD predictions and recurrence predictions differ and different UPOs are converged; (iii) cases where DMD and a recurrence analysis show similar predictions but only the DMD guess converges and finally (iv) cases where DMD yields a converged UPO in the absence of a near recurrence.

3.2.1. Scenario (i): recurrent flow analysis success and DMD failure

There are two instances where both the recurrent flow analysis and our DMD method indicate a likely UPO, but only the guess from the recurrent flow analysis converges in the Newton solver ($t_s \approx 262$ and $t_s \approx 748$ in figure 6). Intriguingly, the guesses for the period generated from the DMD are more accurate than those identified in the recurrence plot. This suggests that the DMD-mode based guesses constructed using (2.11) constitute poor initial conditions for the Newton solver. In instances like these with both a near recurrence and a DMD indicator, a more robust approach could potentially combine elements of both methods, perhaps with a guess that is a concatenation of the state itself and the period coming from DMD, i.e. $[\mathbf{u}(t_i), T_g]^T$.

3.2.2. Scenario (ii): success in both methods but with different UPOs converged

On several occasions both the DMD and recurrent flow analysis identify possible UPOs at the same approximate t_s , but the converged solutions differ. In some of

t_s	T_g	T_ω	% error in T_g	channelflow T_ω	$T_\omega^{RFA}(t_s, \text{ \% error})$
0	67.69	86.43	21.6	—	—
10	62.12	85.60	27.5	85.27	—
55	92.64	95.76	3.3	—	—
115	92.42	97.62	5.3	—	100.17(128,4.2)
175	84.35	77.25	9.1	—	—
180	81.47	76.54	6.4	—	—
200	75.04	78.23	4.1	—	—
235	61.38	71.83	14.5	—	71.83(235,12.3)
410	72.86	95.76	23.9	—	103.01(412,3.7)
415	84.19	104.28	23.9	—	103.01(412,3.7)
545	76.13	86.43	11.9	—	—
715	68.12	64.95	4.8	—	—
725	68.99	94.76	27.2	—	—
730	70.19	64.95	8.0	—	85.60(733,6.5)
735	71.31	88.61	17.3	87.89	85.60(733,6.5)
740	71.58	88.20	19.5	88.90	85.60(733,6.5)
780	88.73	90.71	3.3	90.31	—
815	76.19	78.23	2.6	—	—
825	79.71	70.12	13.7	—	—
880	95.14	104.28	8.8	—	—

TABLE 1. Periodic orbits found from the DMD results reported in figure 6. The guessed period of the UPO, T_g , was determined from $\omega_f(n=2)$ (2.9). In the penultimate column, the UPOs found from DMD-generated guesses are connected to known UPOs converged in the slightly shorter $L_x = 1.75\pi$ box listed on channelflow.org (see also the discussion at the start of § 2). The final column shows a comparison to recurrent flow analysis, T_ω^{RFA} indicates the period of any UPOs found with that method for start times close to that used in the DMD. Note that the UPO found at $t=0$ and $t=545$ (and studied in detail in § 2) has shift–reflect and shift–rotate symmetries (Kawahara 2005).

these cases, it can be argued that one approach is ‘more accurate’ than the other. For example, compare the UPO with period $T_\omega = 103.01$ found from a near recurrence at $t_s = 412$ to the DMD predictions at $t_s \in \{410, 415\}$. As described above, the short DMD time window is unable to accurately parameterise the long UPO with significant errors in the predicted period(s). However, both DMD-based guesses do converge – to UPOs with different periods that are commensurate with the period obtained via recurrent flow analysis. The change in period in the Newton iterations is striking and is not something we have observed in guesses from recurrent flow analysis. There are other times where the UPOs found via DMD and recurrence differ (e.g. $t_s \approx 128$, $t_s \approx 730$) but the differences in the predicted and true periods are less extreme and it is difficult to argue which guess was ‘better’. Here, there is a plausible connection with the project-then-search method described by Ahmed & Sharma (2017) for finding equilibria. In that method, known exact solutions are projected onto a few resolvent modes before being input as new guesses to a Newton solver, resulting in a large number of new equilibria which are qualitatively similar to the known solutions. In our approach, we initialise the Newton solver with a guess constructed from just a few (five for the two-harmonic cases) DMD modes.

These observations raise an interesting ambiguity: given that there appear to be groups of UPOs which are visually very similar (i.e. constitute a closed loop of the

same physical processes) with very similar periods, many of which can be accessed from a Newton search with slightly different initial conditions, can we determine which solution the turbulent flow is actually shadowing? Although we do not explore this here, more insight could potentially be gained by employing a longer observation window (e.g. taken to be slightly longer than any identified near recurrence) and by using more harmonics to generate the initial guess.

3.2.3. Scenarios (iii) and (iv): DMD success with recurrent flow failure, or in the absence of a near recurrence

There is less ambiguity in cases where the recurrent flow analysis guess fails (scenario iii). In these situations the DMD-based approach offers significant advantages. An example of the former scenario is seen in the DMD-based successes for $t_s \in \{815, 825\}$, where two near recurrences are found nearby but neither converges. In both of these cases, the period generated by the DMD-based approach differs non-negligibly from the time T at which $R(t_s, T)$ has a local minima. This is another scenario where the two methods can work well together to converge exact solutions where they may previously have been dismissed.

Perhaps the strongest advantage offered by the DMD approach is in scenario (iv) – a converged UPO from DMD without a corresponding near recurrence; examples of this behaviour are observed at $t_s \in \{180, 200, 780\}$. These results are also notable due to the quality of the guess (e.g. the relative error in the period at $t_s = 780$ is 3.3%). The guess at $t_s = 780$ is particularly interesting since the DMD window corresponds to only 2/3 of the full period, and in figures 7 and 8 we examine it further, comparing the DMD modes used to build the UPO guess to the true Koopman modes of the UPO (the coefficients of a Fourier series). There is good qualitative agreement between the modes, providing further evidence that DMD of modest-duration turbulent trajectories can really see the essence of a nearby exact coherent structure.

It is worth emphasising that this final class of guesses could not be obtained another way as there is no near recurrence in the data, and so this method appears to offer some promise at higher Reynolds numbers where the UPOs tend to be more unstable. A common feature of the various guesses from DMD without a near recurrence is that only the first two harmonics are resolved accurately (see the difference in the bottom two panels of figure 6). This behaviour was also observed for the simple edge state examined in § 2.2, where we observed improved resolution of the higher-frequency Koopman modes as the length of the observation window was increased. However, for turbulent trajectories the observation window cannot be increased arbitrarily and would require a careful adjustment to avoid the inclusion of snapshots where the flow has left the neighbourhood around the UPO where the local Koopman decomposition holds.

3.3. Summary

Overall, the success rates (the ratio of converged solutions to number of guesses) of recurrent flow analysis and of DMD are comparable over the 1000 advective time units, at 0.462 and 0.435 respectively. However, the raw number of UPOs found via DMD is significantly larger than that found by recurrent flow analysis: 20, of which 16 are unique, compared with 6 (all unique).

A caveat in this assessment is that there are different adjustable parameters in each approach, which makes it challenging to compare them quantitatively in detail. For example, the threshold $R(t, T)$ at which a near recurrence is flagged as worth




FIGURE 7. Contours of (the real part of) the streamwise velocity for the first three DMD modes for (a) the time window $T_w = 60$ starting at $t_s = 780$ (see table 1) and (b) the converged UPO with period $T_\omega = 71.83$. In both cases, the number of snapshot pairs is $M = 100$, with $\delta t = 1$. Note modes have been normalised so that contours run -1 (black) to $+1$ (bright yellow).

inputting to the Newton solver could be relaxed, or we could try multiple velocity fields corresponding to times near to where $R(t, T)$ has a local minimum, rather than the minimising t only. In DMD, we could adjust the length of the observation window T_w , the separation time in a snapshot pair δt and the number of snapshots M . However, despite these ambiguities we emphasise that we have not here attempted to optimise our DMD design, and have had success in converging UPOs with a variety DMD parameter choices (not shown). For example, time window lengths of roughly $50 \lesssim T_w \lesssim 100$ can work depending on the period of the nearby UPO. We note in passing the similarity between the length of these windows, the period of many of the UPOs ($T_\omega \sim 100$), and the eddy turnover time scale in this flow ($T \sim 100$, see Hamilton *et al.* (1995)). In addition, the method is fairly insensitive to different values of δt , as long as it is of the order of an advective time unit.

Finally, it is worth comparing the computational complexity of the new DMD approach to that of recurrent flow analysis. For the sake of comparison we compare a single DMD computation on M snapshot pairs with a single recurrence calculation where t is fixed and we search back (or forward) a time T , requiring the comparison of the state at time t with $\sim M$ additional snapshots of the flow. Each snapshot is represented by a vector with $N (= 3 \times N_x \times N_y \times N_z)$ components, and $N \gg M$.

Performing a recurrent flow analysis at a single time instant t requires the computation of an L_2 norm of the difference $\mathbf{u}(t) - \mathbf{u}(t + T)$ over M snapshots.




FIGURE 8. As figure 7, but now showing the streamwise-averaged streamwise velocity (contours) and the streamwise average y - z streamfunction (lines and arrows).

Each subtraction-plus- L_2 norm has a time complexity $O(N)$, so the overall cost when looping over T is $O(MN)$. The DMD approach can be divided into three major steps: (i) an SVD of the data matrix for the pseudo-inverse, with complexity $O(M^2N)$; (ii) an eigendecomposition, which is $O(M^3)$ (assuming that the rank $r \sim M$); and (iii) fitting a handful of modes to the trajectory to build a guess, where the solution of the associated minimisation problem ((2.12), see Page & Kerswell (2019)) has a complexity of $O(MN)$ (assuming only a few modes are used in the fit).

Clearly the expense of the SVD in the DMD renders the approach more computationally expensive than a recurrent flow analysis, although with $M = O(100)$ these computations can still be performed on a laptop. However, note that the recurrent flow analysis presented in figure 6 was performed at a relatively fine temporal resolution, $\delta t = 1$, while DMD computations were only performed every five advective time units. Also, both methods are ultimately generating guesses for a Newton–Krylov search, which is a much more computationally demanding task. The significant increase in the number of UPOs identified with the DMD procedure and its manageable computational cost should warrant its use alongside or in place of a recurrent flow analysis.

4. Discussion

In this section we briefly examine some features of our initial guesses to provide some insight into what constitutes a good guess for UPO. We do this by exploring how well the short-time DMD computations can see Koopman eigenfunctions related

to the structure of interest and by performing DMD on the converged turbulent UPOs themselves to identify a rough rule of thumb on where the best DMD time windows tend to be located in state space.

The idea that DMD can identify the signature of UPOs in a long trajectory is based on the fact that these exact coherent structures possess a local Koopman decomposition which is related to both the period of the solution itself and the eigenvalues governing the dynamics within its linear subspace (Mezic 2019). For short time windows, DMD can typically only resolve a few eigenvalues related to the structure itself (the UPO's Fourier harmonics) as seen in §3 (see also Page & Kerswell (2019)). The 'quality' of a particular DMD-based guess, and some indication as to whether the convergence of a UPO occurred due to luck, can be examined by computing approximations to the Koopman eigenfunctions (see (2.7)) using left eigenvectors from the short-time-window DMD,

$$\varphi_{\lambda_j}^{DMD}(\mathbf{u}) = (\mathbf{w}_j^{DMD})^H \boldsymbol{\psi}(\mathbf{u}), \quad (4.1)$$

and comparing them to those obtained from the converged UPO itself,

$$\varphi_{\lambda_j}(\mathbf{u}) = (\mathbf{w}_j^{UPO})^H \boldsymbol{\psi}(\mathbf{u}). \quad (4.2)$$

The above procedure is performed for two periodic orbits and their original guesses in figure 9, where the velocity field used in the projections is that of the UPO itself. One of the guesses examined was selected due to its particularly poor estimate of the period of the UPO (at $t_s = 10$, $T_g = 62.12$ and $T_\omega = 85.60$, a relative error of 27.5%) and the Koopman eigenfunctions from the windowed DMD (4.1) compare poorly with the 'true' Koopman eigenfunctions evaluated on the UPO itself (4.2). Furthermore, the projection shows sharp gradients and appears almost discontinuous in places. On the other hand, the second guess was considerably more accurate and the approximation to the Koopman eigenfunctions is qualitatively much better. For other UPOs we have observed similar behaviour (not shown); that more accurate initial guesses (as measured by the predicted period) tend to produce Koopman eigenfunctions which resemble qualitatively those of the UPO itself.

Computation of the Koopman eigenfunctions along the UPO itself is only possible after the solution has been converged, and it is natural to speculate whether there are additional indicators of guess quality that can be assessed before the Newton solver is initialised. While the eigenspectrum-based measure ε_ω (2.10) provides a relatively robust classifier to identify sensible initial guesses, we have observed some additional heuristics that can be used to screen a large number of initial guesses. When examining a long time series with short DMD windows, perhaps the most obvious clue is the observation of 'clusters' of similar guesses (in terms of period T_g) in time. This is evident in some of the successful guesses highlighted in figure 6 and could be further emphasised by shrinking the timestep between subsequent DMD calculations (in figure 6 this is $\Delta_w = 5$).

Another interesting feature of good guesses from short DMD windows is that they tend to share certain characteristics in terms of their location in an energy production–dissipation plane. To demonstrate this effect, we perform many DMD calculations with fixed window length $T_w = 60$ on two of the longer converged UPOs and plot the 'best' and 'worst' time windows in figure 10, where the criterion used to rank the quality of the DMD spectra was the error in the guessed period (using the first two harmonics),




FIGURE 9. Real and imaginary components of approximate Koopman eigenfunctions, $\varphi = \varphi^r + i\varphi^i$, corresponding to the eigenvalues $\lambda_1 \approx i\omega_f$ and $\lambda_2 \approx 2i\omega_f$, obtained in DMD applied to a converged UPO (solid lines) and the fixed turbulent window $T_w = 60$ that converged to that same UPO (dashed lines). In both cases the eigenfunction is evaluated along the UPO itself. (a,b) The initial guess at $t_s = 10$ (see table 1) and the converged UPO $T_\omega = 85.60$. (c,d) The initial guess at $t_s = 235$ and the converged UPO $T_\omega = 71.83$.

$|2\pi/\omega_f(n = 2) - T_\omega|/T_\omega$. The best guesses tend to sample the upper-right quadrant (the fast, high production, high dissipation region) at the expense of the slower, gentler region in the lower-left corner. Note that for longer orbits, good estimates (relative error $\lesssim 5\%$) can also be obtained on trajectories which mostly shadow the portion of the curve $D < I$ (not shown) – the key requirement appears to be the inclusion of the most extreme values of production and dissipation.

We note that the presentation in this manuscript has focused on one particular time window length and DMD design in order to introduce the method in a clear and systematic way. There is clearly scope for the results to be improved by adjusting the DMD design based on local performance. For example, where near recurrences are observed it would make sense to increase the time window to the predicted period, while for cases where no near recurrence is observed but DMD does predict a UPO, it would be interesting to search for a ‘best’ local guess by small adjustments in window length and position.

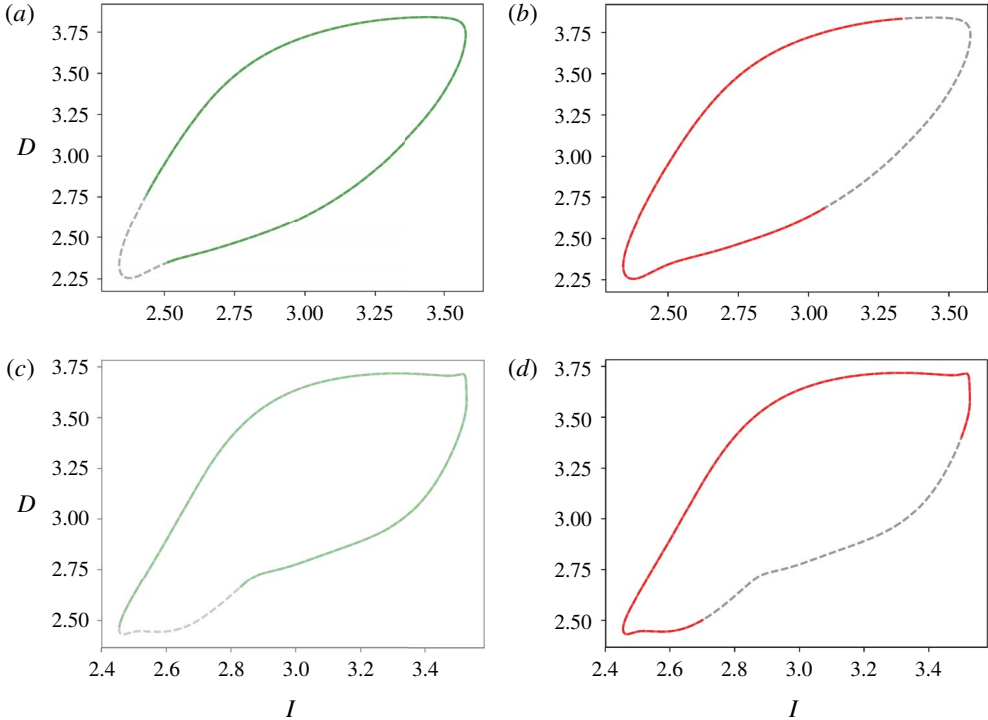


FIGURE 10. ‘Best’ (left, in green) and ‘worst’ (right, in red) locations for a DMD time window of length $T_w = 60$ as measured by the difference in predicted period from that of the true UPO, $|T_g - T_\omega|/T_\omega$, visualised in a two-dimensional projection of energy production against dissipation (both normalised by their values in laminar flow). The computations are performed on the UPOs themselves, and the full UPO is shown in dashed grey. The predicted period $T_g = 2\pi/\omega_f$, where ω_f was computed from the first two harmonics. (a,b) Background UPO has period $T_\omega = 85.60$. (c,d) Background UPO has period $T_\omega = 94.76$.

5. Conclusion

We have presented a new method based on DMD to both identify the signature of nearby UPOs in time series of turbulent flows, and to generate robust initial guesses that can serve as inputs to a Newton–Krylov algorithm. The approach is designed to be applied to short turbulent trajectories, with the existence of a nearby UPO deemed ‘likely’ when the DMD eigenvalue spectrum exhibits repeated harmonics of a fundamental frequency. Initial conditions for the Newton solver are then built from these near-periodic DMD modes and a guess for the period based on the estimated fundamental frequency.

The ability of the method to function without the need for a near recurrence within the observation window was demonstrated on a simple UPO before we applied it to a long turbulent trajectory. With a modest time window (shorter than most known UPOs), the new method identified many more UPOs than were found with a recurrent flow analysis of the same dataset. These solutions were found in regions where (i) a near recurrence was found and converged, (ii) a near recurrence was identified but the Newton solver failed to converge and (iii) where no near recurrence was flagged. These outcomes indicate that the method can be both a

useful complementary tool and a potentially powerful alternative to recurrent flow analysis. In particular, the success of the method in regions where no near recurrence occurs suggests it may be of use at higher Reynolds numbers where the probability of shadowing a UPO for a full cycle drops.

An interesting behaviour was found at times where both the DMD based method and recurrent flow analysis flagged the probable existence of a UPO. In some cases, the DMD approach converged to a different solution than the guess from a recurrent flow analysis. Often, the error in the predicted period for these alternate solutions was comparable to the recurrent flow analysis, and given the sensitivity of the Newton method to initial conditions it is difficult to assert which one was indeed the solution being shadowed in phase space. Motivated by the variable discrepancy between the predicted period and that of the converged solution, we also presented some analysis of the features of both ‘good’ and ‘bad’ initial guesses generated by the method. Good guesses tend to require the DMD to see the most extreme dissipation events along an orbit at the expense of the slower, more gentle dynamics.

There are a number of interesting avenues open to further investigation. The most natural question is whether the method can be extended to much higher Reynolds numbers to converge UPOs, where recurrent flow analysis struggles to identify any guesses at all (e.g. see Chandler & Kerswell (2013) and Lucas & Kerswell (2015)). There are also some intriguing properties of the DMD-based approach that merit further study. For example, DMD is known to struggle with resolving a mixture of time periodicity and translation which characterises relative periodic orbits (RPOs) as each frequency in the translating frame is split into many in the fixed frame. However, DMD can still be effective for RPOs when the data are preprocessed, or ‘pulled back’ to a particular reference frame (Willis *et al.* 2013) by fixing its phase relative to some template velocity fields (N. Budandur; private communication). Another issue is that DMD on short time windows performs ‘best’ (i.e. the predicted period matches that of the underlying simple invariant set) when it sees the faster, high-dissipation region of the UPO. In instances where the observation window does not have this property, it would be interesting to see if the performance of the DMD could be improved, for instance by modifying how snapshot pairs are distributed in time.

Acknowledgement

We are very grateful to Y. Dubief for sharing his viscoelastic code with us.

Declaration of interests

The authors report no conflict of interest.

REFERENCES

- AHMED, M. A. & SHARMA, A. S. 2017 New equilibrium solution branches of plane Couette flow discovered using a project-then-search method. [arXiv:1706.05312](https://arxiv.org/abs/1706.05312).
- ARBABI, H. & MEZIĆ, I. 2017 Study of dynamics in post-transient flows using Koopman mode decomposition. *Phys. Rev. Fluids* **2**, 124402.
- BAGHERI, S. 2013 Koopman-mode decomposition of the cylinder wake. *J. Fluid Mech.* **726**, 596–623.
- BRUNTON, S. L., BRUNTON, B. W., PROCTOR, J. L. & KUTZ, J. N. 2016 Koopman invariant subspaces and finite linear representations of nonlinear dynamical systems for control. *PLoS ONE* **11** (2), e0150171.

- CHANDLER, G. J. & KERSWELL, R. R. 2013 Invariant recurrent solutions embedded in a turbulent two-dimensional Kolmogorov flow. *J. Fluid Mech.* **722**, 554–595.
- CVITANOVIĆ, P., ARTUSO, R., MAINIERI, R., TANNER, G. & VATTAY, G. 2016 *Chaos: Classical and Quantum*. Niels Bohr Inst.
- CVITANOVIĆ, P. & GIBSON, J. F. 2010 Geometry of the turbulence in wall-bounded shear flows: periodic orbits. *Phys. Scr.* **T142**, 014007.
- DUBIEF, Y., TERRAPON, V. E., WHITE, C. M., SHAQFEH, E. S. G., MOIN, P. & LELE, S. K. 2005 New answers on the interaction between polymers and vortices in turbulent flows. *Flow Turbul. Combust.* **74** (4), 311–329.
- HALL, P. & SHERWIN, S. 2010 Streamwise vortices in shear flows: harbingers of transition and the skeleton of coherent structures. *J. Fluid Mech.* **661**, 178–205.
- HAMILTON, J. M., KIM, J. & WALEFFE, F. 1995 Regeneration mechanisms of near-wall turbulence structures. *J. Fluid Mech.* **287**, 317–348.
- JOVANOVIĆ, M. R., SCHMID, P. J. & NICHOLS, J. W. 2014 Sparsity-promoting dynamic mode decomposition. *Phys. Fluids* **26**, 024103.
- KAWAHARA, G. 2005 Laminarization of minimal plane Couette flow: going beyond the basin of attraction of turbulence. *Phys. Fluids* **17**, 041702.
- KAWAHARA, G. & KIDA, S. 2001 Periodic motion embedded in plane couette turbulence: regeneration cycle and burst. *J. Fluid Mech.* **449**, 291–300.
- KAWAHARA, G., UHLMANN, M. & VAN VEEN, L. 2012 The significance of simple invariant solutions in turbulent flows. *Annu. Rev. Fluid Mech.* **44** (1), 203–225.
- KERSWELL, R. R. 2005 Recent progress in understanding the transition to turbulence in a pipe. *Nonlinearity* **18**, R17–R44.
- KOOPMAN, B. O. 1931 Hamiltonian systems and transformations in Hilbert space. *Proc. Natl Acad. Sci. USA* **17** (5), 315–318.
- KUTZ, J. N., BRUNTON, S. L., BRUNTON, B. W. & PROCTOR, J. L. 2016 *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*, 1st edn. SIAM.
- LUCAS, D., CAULFIELD, C. P. & KERSWELL, R. R. 2017 Layer formation in horizontally forced stratified turbulence: connecting exact coherent structures to linear instabilities. *J. Fluid Mech.* **832**, 409–437.
- LUCAS, D. & KERSWELL, R. R. 2015 Recurrent flow analysis in spatiotemporally chaotic 2-dimensional Kolmogorov flow. *Phys. Fluids* **27**, 045106.
- MEZIĆ, I. 2005 Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dyn.* **41**, 309–325.
- MEZIĆ, I. 2013 Analysis of fluid flows via spectral properties of the Koopman operator. *Annu. Rev. Fluid Mech.* **45**, 357–378.
- MEZIC, I. 2019 Spectrum of the Koopman operator, spectral expansions in functional spaces, and state-space geometry. *J. Nonlinear Sci.* doi:10.1007/s00332-019-09598-5.
- MEZIC, I. & BANASZUK, A. 2004 Comparison of systems with complex behavior. *Physica D* **197**, 101–133.
- PAGE, J. & KERSWELL, R. R. 2018 Koopman analysis of Burgers equation. *Phys. Rev. Fluids* **3**, 071901(R).
- PAGE, J. & KERSWELL, R. R. 2019 Koopman mode expansions between simple invariant solutions. *J. Fluid Mech.* **879**, 1–27.
- ROWLEY, C. W. & DAWSON, S. T. M. 2017 Model reduction for flow analysis and control. *Annu. Rev. Fluid Mech.* **49**, 387–417.
- ROWLEY, C. W., MEZIĆ, I., BAGHERI, S., SCHLATTER, P. & HENNINGSON, D. S. 2009 Spectral analysis of nonlinear flows. *J. Fluid Mech.* **641**, 115–127.
- SCHMID, P. J. 2010 Dynamic mode decomposition of numerical and experimental data. *J. Fluid Mech.* **656**, 5–28.
- SCHMID, P. J., LI, L., JUNIPER, M. P. & PUST, O. 2011 Applications of the dynamic mode decomposition. *Theor. Comput. Fluid Dyn.* **25**, 249–259.
- SCHNEIDER, T. M., ECKHARDT, B. & YORKE, J. A. 2007 Turbulence transition and the edge of chaos in pipe flow. *Phys. Rev. Lett.* **99**, 034502.

- TU, J. H., ROWLEY, C. W., LUCHTENBURG, D. M., BRUNTON, S. L. & KUTZ, J. N. 2014 On dynamic mode decomposition: theory and applications. *J. Comput. Dyn.* **1** (2), 391–421.
- VISWANATH, D. 2007 Recurrent motions within plane Couette turbulence. *J. Fluid Mech.* **580**, 339–358.
- WALEFFE, F. 1997 On a self-sustaining process in shear flows. *Phys. Fluids* **9**, 883–900.
- WANG, J., GIBSON, J. & WALEFFE, F. 2007 Lower branch coherent states in shear flows: transition and control. *Phys. Rev. Lett.* **98**, 204501.
- WILLIAMS, M. O., KEVREKIDIS, I. G. & ROWLEY, C. W. 2015 A data-driven approximation of the Koopman operator: extending dynamic mode decomposition. *J. Nonlinear Sci.* **25** (6), 1307–1346.
- WILLIS, A. P., CVITANOVIC, P. & AVILA, M. 2013 Revealing the state space of turbulent pipe flow by symmetry reduction. *J. Fluid Mech.* **721**, 514–540.