



Intrinsic phase-based proper orthogonal decomposition (IPhAB POD): a method for physically interpretable modes in near-periodic systems

Akhileshwar Borra^{1,†}, Zoey Flynn¹, Andres Goza¹ and Theresa Saxton-Fox¹

¹Department of Aerospace Engineering, University of Illinois, Urbana, IL 61801, USA

(Received 17 April 2024; revised 23 September 2024; accepted 18 October 2024)

Fluid dynamics systems driven by dominant, near-periodic dynamics are common across wakes, jets, rotating machinery and high-speed flows. Traditional modal decomposition techniques have been used to gain insight into these flows, but can require many modes to represent key physical processes. With the aim of generating modes that intuitively convey the underlying physical mechanisms, we propose an intrinsic phase-based proper orthogonal decomposition (IPhAB POD) method, which creates energetically ranked modes that evolve along a characteristic cycle of the dominant near-periodic dynamics. Our proposed formulation is set in the time domain, which is particularly useful in cases where the cyclical content is imperfectly periodic. We formally derive IPhAB POD within a POD framework that therefore inherits the energetically ranked decomposition and optimal low-rank representation inherent to POD. As part of this derivation, a dynamical systems representation is utilized, facilitating a definition of phase within the system's near-periodic cycle in the time domain. An expectation operator and inner product are also constructed relative to this definition of phase in a manner that allows for the various cycles within the data to demonstrate imperfect periodicity. The formulation is tested on two sample problems: a simple, low Reynolds number aerofoil wake, and a complex, high-speed pulsating shock wave problem. The method is compared to space-only POD, spectral POD (SPOD) and cyclostationary SPOD. The method is shown to better isolate the dominant, near-periodic global dynamics in a time-varying IPhAB mean, and isolate the tethered, local-in-phase dynamics in a series of time-varying modes.

Key words: shock waves, low-dimensional models, wakes

† Email address for correspondence: akhil.borra@gmail.com

1. Introduction

Dominant, near-periodic dynamics is common across many flows, including bluff body wakes, rotating machinery, and shock-laden processes. These systems often exhibit near-periodic global dynamics, as well as turbulent, smaller-scale behaviours that are tethered to the larger-scale flow. Extracting the disparate flow scales and structures and their interactions is critical to understanding, modelling and controlling these systems. We focus in this paper on developing a data-driven decomposition that provides more physically interpretable representations, or ‘modes’, of these near-periodic systems. (See e.g. Taira *et al.* (2017) for a review of data-driven and operator-based decomposition techniques, which include linear stability analysis, resolvent analysis, proper orthogonal decomposition (POD), and dynamic mode decomposition.) When used to analyse fluid flows generally, most modal decomposition techniques are not guaranteed to provide insight into instantaneously relevant coherent structures. However, in many cases the resulting modes have been used successfully to investigate the physics of specific scales of motion. Comparing modes between forced and unforced cases can identify the impact of forcing on different flow scales in a jet (Feng, Wang & Pan 2011; Schmid *et al.* 2011; Kourentis & Konstantinidis 2012), or can clarify the nature of coherent structures associated with upstream and downstream aero-optic distortions in a turbulent boundary layer (Borra & Saxton-Fox 2022). Modes associated with specific frequencies can isolate the coherent structures that generate bands of acoustic noise (Freund & Colonius 2009; Riberio & Wolf 2017) and those associated with specific instabilities (Huang *et al.* 2016; Lengani *et al.* 2017; Schmidt *et al.* 2018; Pickering *et al.* 2020). Decompositions have advanced understanding of the physics of bluff body wakes (Feng *et al.* 2011; Kourentis & Konstantinidis 2012; Zhang, Lie & Wang 2014; Riberio & Wolf 2017; Wu *et al.* 2017; Ping *et al.* 2021), flows past flexible bodies (Schmid 2010; Goza & Colonius 2018), cavity flows (Murray, Sällström & Ukeiley 2009; Seena & Sung 2011; Hossain, Bergstrom & Chen 2015; Desikan *et al.* 2022) and shear layers (Hilberg, Lazik & Fiedler 1994; Szubert *et al.* 2015; Weightman *et al.* 2018). With improved understanding of the behaviour of specific flow features comes an improved ability to control the flow (Rowley & Dawson 2017). Reduced-order models constructed through modal analysis techniques have been used to design control algorithms for myriad applications, including recirculation control in channel flow (Ravindran 2000), separation control over a backwards facing ramp (Taylor & Glauser 2004), a flat plate at large angles of attack (Ahuja & Rowley 2010), and the wake of a circular cylinder (Bergmann, Cordier & Brancher 2005).

Despite the utility of these decomposition techniques, we highlight three limitations that have motivated modifications within the community. First, while in some flows modes derived from decompositions have been found to physically resemble instantaneous structures (Saxton-Fox & McKeon 2017), in other flows the modes generated by these decompositions are structurally dissimilar from instantaneous flow structures. To understand coherent structure behaviour, researchers have added multiple modes together (Riberio & Wolf 2017) or turned to other techniques entirely, such as finite time Lyapunov exponent fields (Kourentis & Konstantinidis 2012) or conditional projection averaging (Saxton-Fox, Lozano-Durán & McKeon 2022). Second, modal decompositions tend to prioritize behaviours that are energetically or dynamically significant statistically, but these modes can miss important rare or extreme events (Schmidt & Schmid 2019). The modes that are identified as most dynamically or energetically significant can vary depending on the length of the time series considered in multi-scale systems (Mendez, Balabane & Buchlin 2019). Finally, many of the decomposition techniques are appropriate only for statistically stationary flows.

Modifications to standard variants of modal decomposition techniques have been proposed to address some of these challenges. Non-stationarity has been addressed through space–time variants of POD by Gordeyev & Thomas (2013) for use with control, and by Schmidt & Schmid (2019) to identify mechanisms driving rare loud events in jet noise. Multi-resolution techniques are also naturally suited to non-stationary problems, and multi-resolution dynamic mode decomposition, multi-scale POD, and data-driven wavelet techniques have been demonstrated to provide insights into optimal modes over a range of time scales (Kutz, Fu & Brunton 2016; Mendez *et al.* 2019; Floryan & Graham 2021).

For the systems of focus in this paper, which are dominated by near-periodic behaviours, Padovan, Otto & Rowley (2020) and Padovan & Rowley (2022) extend the traditional assumption of a stationary base flow to include periodic base flows for resolvent analysis. This extension enables an improved understanding of how perturbations in a periodic base flow interact with one another and the base flow itself, elucidating cross-frequency modal interactions. Phase-conditioned localized spectral proper orthogonal decomposition (PCL-SPOD) (Franceschini *et al.* 2022) and cyclostationary SPOD (CSPOD) (Heidt & Colonius 2024) both consider POD within a phase-locked, periodic system. Franceschini *et al.* (2022) leverage a windowed SPOD technique, using short-time Fourier transforms, to study the SPOD modes of small-scale turbulent activity within a larger phase variation of a periodic system, e.g. vortex shedding. They also develop a quasi-steady approach that allows this technique to be computed more simply when the large-scale frequencies are much lower than the small-scale frequencies. Heidt & Colonius (2024) also extend SPOD to globally periodic systems, and identify which spectral modes most strongly interact. The approach offers a way of identifying how modes at different frequencies interact with each other and an underlying periodic base state. Heidt & Colonius (2024) also demonstrate that CSPOD is the data-driven analogue to the harmonic resolvent formulation of Padovan *et al.* (2020). These techniques, which focus on decompositions around intrinsic near-periodic dynamics, also share commonalities to the work of Zhang *et al.* (2023), who utilized phase averaging around a periodic forcing to understand modal response linked to external periodic stimuli.

We aim to build on existing work by introducing a new formulation, intrinsic phase-based POD (IPhB POD), that enables near-periodic dynamics to be explicitly accounted for in the decomposition. Instead of requiring a known form of periodic dynamics, IPhB POD utilizes a phase portrait from a user-provided low-dimensional proxy for the global behaviour. Characteristic periodic behaviour is defined from this dynamical systems representation, and cycles from the data that fall within this representative behaviour are selected systematically. A POD formulation is constructed for these cycles that are classified as near-periodic, within the time domain. Working within the time domain avoids the need to represent the system evolution at a single frequency, since for many flows of interest, the evolution is not exactly periodic. The provided time domain formulation uses a clear notion of the phase of the system, defined as a fraction of the trajectory along the full cycle in the phase portrait representation. This phase portrait approach allows for cycles with disparate time periods or large-scale trajectories to be compared meaningfully, provided that they correspond to the same near-periodic behaviour of interest.

As a consequence of the time domain formulation, the modes generated from IPhB POD evolve in time rather than in the frequency domain. Utilizing the time domain could increase representation efficiency and avoid a more broadband representation in frequency space when the underlying dynamics (and the modes connected to it) is not

perfectly periodic. This way of defining the phase of the near-periodic dynamics may be of interest to frequency-based techniques as well. We will later draw connections between IPHaB POD and conditional space–time POD (Schmidt & Schmid 2019), where the event condition here is a near-periodic dynamic cycle that is intrinsic to the system. We show below that this condition requires substantive changes to the POD formulation.

In this paper, we apply IPHaB POD to two examples: a flat plate experiencing vortex shedding (Towne *et al.* 2023), and a periodically moving shock front over a cone-cylinder (Sasidharan & Duvvuri 2021). For both examples, the dominant, near-periodic dynamics is captured by an IPHaB mean, which is a phase average that leverages our unique definition of intrinsic phase. For the unsteady shock front data, in which there are additional, complex dynamics, the IPHaB modes are able to represent local-in-phase dynamics, including supersonic jets and turbulence within the shock structure, that is not visible in modes from space-only POD, SPOD or CSPOD.

2. Data sets

To develop and demonstrate the method discussed in this paper, two data sets are used. The first is from a high-fidelity computation of an incompressible low Reynolds number flow with only one dominant coherent feature that is periodic (to within discretization errors). The second is from an experiment of a high Reynolds number and Mach number flow that has multiple dynamic behaviours interacting with one another. The first data set is used to introduce the method and its properties, and the second to demonstrate the promise of the method on more complex data.

The low Reynolds number case is of flow past a flat plate at angle of attack $\alpha = 30^\circ$, with Reynolds number based on the plate length $Re = 100$, and is acquired from the database for reduced-complexity modelling of fluid flows (Towne *et al.* 2023). At this high (stalled) angle of attack, the long-time behaviour after an initial transient is a globally stable limit cycle attractor in which there is alternating vortex shedding from either side of the body. This periodic limit cycle is shown in figures 1(a–d), which show vorticity at four phases in the periodic shedding cycle. The solution for this case was computed using the immersed boundary method of Colonius & Taira (2008) and Taira & Colonius (2007). The simulation parameters are given non-dimensionally as (normalized using the plate length and free-stream flow speed) $\Delta x = 0.02$ and $\Delta t = 0.1$. The simulation was run for a dimensionless time $t = 50$ prior to data collection to remove the effects of the transient startup. Data were collected for a dimensionless time $t = 40$, or 9 vortex-shedding cycles.

The data set involving more complex dynamics, used to demonstrate the promise of the method, comes from the schlieren experiments by Sasidharan & Duvvuri (2021) of high-speed flow past a cone-cylinder body. The dynamics of the flow field is controlled by three non-dimensional geometric parameters: the Mach number (M), the ratio L/D of the cone length (L) to the diameter of the cylinder base (D), and the half-angle of the cone (θ). The parameters chosen for this study are $M = 6$, $L/D = 0.5$ and $\theta = 25^\circ$. For the IPHaB POD analysis, we used time-resolved schlieren courtesy of Sasidharan & Duvvuri (2021), with sampling frequency 57 kHz. These data were collected for 0.0983 s. For these parameters, large-scale pulsations in the shock location are observed. Four instances of the shock pulsations are shown in figures 1(e–h), highlighting the various shocks and their interactions, as well as key processes tethered to that shock interplay (boundary layer development and separation, and the supersonic jet near the triple point). The data were of high quality but naturally included both high- and low-frequency noise. The low-frequency noise was characterized by a slow, monotonic change in the intensity of the light in the

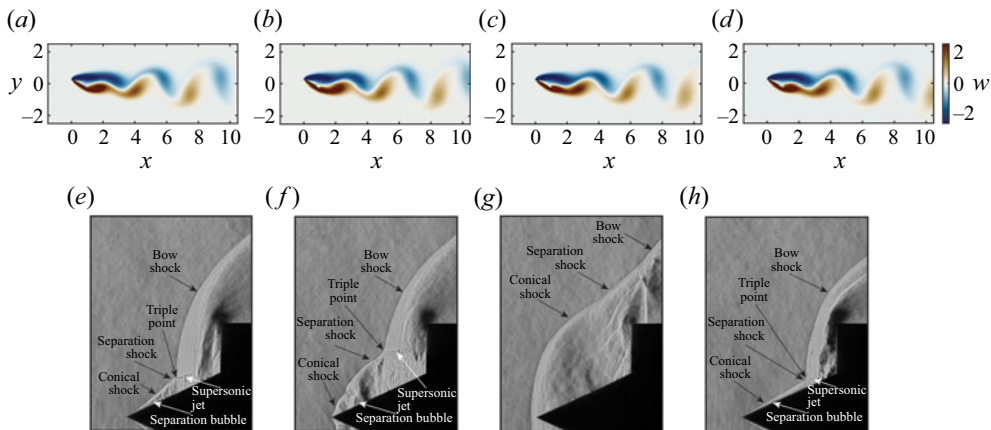


Figure 1. Four instances of (a–d) vortex shedding from flat plate at a stalled angle of attack $\alpha = 30^\circ$ (Towne *et al.* 2023), and (e–h) shock pulsations over a cone-cylinder body (Sasidharan & Duvvuri 2021). Vorticity from a two-dimensional computation is plotted in (a–d). A sample movie of the vorticity field is available in supplementary movie 1 (available at <https://doi.org/10.1017/jfm.2024.1061>). Experimental schlieren data are shown in (e–h). A sample movie of the schlieren imaging over four shock pulsation cycles is available in supplementary movie 2.

overall image as a function of time, while the high-frequency noise was characterized by rapid changes in the pixel intensity. The low-frequency noise was mitigated prior to performing the analysis by normalizing the pixel intensity by the local-in-time spatial average of the ‘free-stream’ intensity. For this part of the treatment, this free-stream was defined as the region of the flow into which the shock never penetrated across all time instances. The high-frequency noise was not mitigated prior to implementing the data analysis method described in this work.

We provide a brief summary of the flow physics resulting in the shock pulsations here, for context, but more information regarding the data and the system’s physics can be found in Sasidharan & Duvvuri (2021). The shock pulsations arise due to complex interactions between a bow shock and conical shock. The intersection of the conical and bow shocks results in a transmitted shock that impinges on the boundary layer flow on the cone surface, generating a separation bubble and separation shock with a shock angle larger than the conical shock. The shock front consisting of the conical, separation and bow shocks will be called the shock system. Downstream of the transmitted shock, a supersonic jet forms as seen in figures 1(e,f,h). The size of the separation bubble increases, and results in the separation point moving upstream on the cone surface, the separation shock angle growing, and the triple point moving upwards (cf. figure 1f). Eventually, the entire cone-cylinder body is encompassed by the shock system (cf. figure 1g). The separation bubble then collapses, the shock system rapidly moves closer to the base cylinder, and the cycle restarts (cf. figure 1h). This unsteady process is near-periodic, i.e. there is a pattern of shock pulsations that repeats cyclically (the shock motion), but with imperfect periodicity because multiple important scales yield variations in the behaviour across each cycle.

3. Formulation for IPHaB POD

The proposed formulation aims to decompose data that include a global, nearly periodic cycle. We emphasize that we allow for the overall dynamical system to be imperfect in its periodicity. The (near-)periodic dynamics will be assumed to have a nominal period

duration \bar{T} , with allowed deviations from this nominal trajectory. Because of a lack of perfect periodicity, a naive attempt to define phase as a fixed number of time steps along some k th cycle of the system is inappropriate. Inevitably, the second sample from cycle k will align with a different phase of the near-periodic dynamics from the second sample from cycle m unless care is taken to ensure that they are taken at the same phase of the near-periodic dynamics. This challenge was observed by Franceschini *et al.* (2022): in performing their POD formulation to analyse how a bluff body wake modified turbulent small scales, the authors identified a lack of periodicity in the vortex shedding as an inhibitor to computing meaningful modes. Even for perfectly periodic flows, the need for finite sampling in time can cause a similar outcome: unless the signal is sampled at exactly the same instances along a period (or with sufficient resolution that the same instances along the cycle can be extracted), the same spurious outcome of having modes inherit (an artificially introduced) non-periodicity will occur.

We propose a POD formalism that directly acknowledges the imperfect periodicity of the dynamics while retaining an underlying near-periodic character. Within this setting, we define an inner product that matches each realization to a common phase of a nominal (near-)periodic orbit. The results demonstrate an improved ability to isolate the near-periodic behaviour from the dynamics that is driven by it. In this section, we characterize the form assumed for the near-periodic dynamics, define an expectation operator and inner product informed by this definition of near-periodicity, and describe the implementation of the resulting algorithm.

3.1. Characterizing the quasi-periodic dynamics

To probe the quasi-periodic dynamic behaviour, we work within a spatially discrete formulation of the governing equations, which we write generically as

$$\mathbf{M}\dot{\mathbf{q}} = \mathbf{f}(\mathbf{q}). \quad (3.1)$$

In (3.1), \mathbf{q} conveys the flow state sampled at a subset of points n_s within the flow domain Ω ; e.g. in an incompressible setting, \mathbf{q} is a concatenation of the spatially discrete flow state variables. (For an incompressible flow, these would be the various velocity components and pressure at the collection of discrete spatial points used to represent Ω .) The vector \mathbf{f} is the dynamical evolution operator. The matrix \mathbf{M} allows for non-unity weightings on the time derivatives, including a singular \mathbf{M} as is needed to represent conservation of mass in the incompressible Navier–Stokes equations (resulting in a differential–algebraic system of Hessenberg index 2).

We assume that the dynamics (3.1) contains underlying near-periodic dynamic behaviour, and that the user has a means of identifying distinct cycles of this behaviour. That is, IPHaB POD requires that an operation \mathcal{F} be known such that $\mathbf{q}_l = \mathcal{F}(\mathbf{q})$ provides a meaningful way to probe the intrinsic near-periodic content (without the additional dynamics driven by this near-periodic content). That is, \mathbf{q}_l exhibits the targeted near-periodic dynamics that allows for distinct cycles, and the phase along each cycle, to be identified. We refer to the dynamics of \mathbf{q}_l (i.e. the dynamics of \mathbf{q} under \mathcal{F}) as restricted dynamics. While the requirement of an appropriate \mathcal{F} may appear somewhat limiting, in many cases such knowledge is available. Examples of flows with identifiable near-periodic content include wakes, separation bubble ‘breathing’, fluid–structure interaction problems, and rotating systems. We will next characterize mathematically the assumed near-periodic behaviour, and give examples of this mathematical representation for the two representative cases described in § 2.

We assume that the restricted dynamics $\mathbf{q}_l = \mathcal{F}(\mathbf{q})$ yields n_c cycles. We further assume that the restricted trajectory associated with the k th cycle, $\{\mathbf{q}_l(t) : t \in [t_{0_k}, t_{f_k}]\}$, $k = 1, 2, \dots, n_c$, has a (near-)period $T_k := t_{f_k} - t_{0_k}$. Associated with each k th cycle is also a minimum and maximum at, say, t_{min_k} and t_{max_k} , given by derivative conditions of \mathbf{q}_l over that cycle. The peak-to-peak amplitude for each cycle can thus be defined as $A_k := \|\mathbf{q}_l(t_{max_k}) - \mathbf{q}_l(t_{min_k})\|$. One can define a nominal period duration and peak-to-peak amplitude from the arithmetic mean of the quantities T_k and A_k over all cycles, e.g. $\bar{A} = (1/n_c) \sum_{k=1}^{n_c} A_k$ (with an analogous definition for \bar{T}).

Because the dynamics is assumed to be imperfectly periodic, in general $\mathbf{q}_l(t_{f_k}) \neq \mathbf{q}_l(t_{0_k})$. In fact, there may be cycles for which the period duration is very different from the nominal period \bar{T} , and the restricted trajectory $\{\mathbf{q}_l(t) : t \in [t_{0_k}, t_{f_k}]\}$ yields large values of $\|\mathbf{q}_l(t_{f_k}) - \mathbf{q}_l(t_{0_k})\|$. For the purposes of this paper, we focus on restricted trajectories that exhibit near-periodic orbits, and omit extreme events from consideration. That is, we perform IPHaB POD only on cycles that are near-periodic in the sense that their restricted trajectories in phase space are closed, and their associated time durations are equal to a representative period duration (to within some tolerance). The near-periodic trajectories considered within IPHaB POD are then computed as those satisfying

$$\gamma_j = \{\mathbf{q}_l(t) : t \in [t_{0_j}, t_{f_j}], |\bar{T} - T_j|/\bar{T} := \delta_j < \delta^*, \|\mathbf{q}_l(t_{f_j}) - \mathbf{q}_l(t_{0_j})\|/\bar{A} := \epsilon_j < \epsilon^*\}, \quad (3.2)$$

where δ^* and ϵ^* are threshold values for departure from the nominal period duration and from trajectory closedness. We denote the number of near-periodic orbits satisfying (3.2) as n_p , the set of indices associated with these orbits as N (with $\dim(N) = n_p \leq n_c$), and $\Gamma = \{\gamma_j, j \in N$.

It will be useful in presenting IPHaB POD to define the ‘most’ representative periodic cycle of the restricted dynamics, which we denote as the trajectory $\gamma^p := \{\mathbf{q}_l(t) : t \in [t_{0_k}, t_{f_k}]\}$, $k \in N$, that satisfies $\min_k \sqrt{\delta_k^2 + \epsilon_k^2}$. (Note that by construction, γ^p is one of the admitted near-periodic trajectories in (3.2).) We denote the period and discrepancy from trajectory closedness of the nominally periodic γ^p as T^p and ϵ^p , respectively.

We now discuss how the trajectories $\gamma_j, j \in N$, and γ^p are obtained for the two problems considered in this paper. In addition to illustrating how these trajectories are determined, the examples are also intended to motivate as modestly reasonable the assumption that a means to probe the near-periodic content within the dynamics, given by a user-provided mapping \mathcal{F} , is available.

For the aerodynamic flat plate problem, vortex shedding is the (only) content within the dynamics, and is perfectly periodic to within numerical error (cf. figures 1(a–d) and the surrounding discussion). The lift dynamics can be used to characterize this vortex-shedding behaviour. That is, while the lift is not equivalent to the large-scale vortex shedding, its dynamics can be used to identify distinct phases of this vortex shedding (i.e. periodic) behaviour. Thus, for this problem, an \mathcal{F} can be defined that takes the flow state, extracts the flow stresses via the gradient of the pressure field and the viscous Laplacian of the flow velocity, and evaluates these on the surface of the aerodynamic body. (Or alternatively, many immersed boundary methods provide the surface stresses directly as a state variable, and under this formulation the surface stress state variable may be spatially integrated.)

The periodic trajectories Γ can be extracted from this definition of \mathcal{F} as follows (as illustrated in figure 2). The lift signal depicted in figure 2(a), plotted sufficiently far after the initial transient that the limit cycle has been reached, demonstrates the periodic behaviour of the system. The variation in trajectory closedness of the lift cycle and




Figure 2. (a) The coefficient of lift C_l for the representative cycle of vortex shedding from the flat plate is shown in red, and the remaining ensembles are plotted in green. (b) Nearness of cycle start and end positions (ϵ) and difference of cycle period to nominal period (δ) for all cycles. Threshold values used to determine near-periodic cycles are shown as dashed lines (for this case, all cycles are considered nearly periodic). (c) Cycles plotted using a phase portrait with the coefficient of lift plotted against its time derivative. All cycles lie on top of each other because all are nearly periodic.

discrepancy from the nominal period duration are shown in figure 2(b). For this problem, the cycle-to-cycle differences are small and due to discretization errors, so thresholds $\epsilon^* = 5 \times 10^{-5}$ and $\delta^* = 2 \times 10^{-3}$ are chosen to include all cycles. For illustration purposes, and to indicate a phase trajectory undergone by the lift as a representation periodic dynamics, figure 2(c) shows the various lift cycles superimposed on one another as a phase portrait of dC_l/dt versus C_l , though only a single curve is visible because of how nearly periodic the system is. The cycle depicted in red in figures 2(a,c) is the one computed as γ^p based on the criterion described above, though again the periodicity of the system dynamics renders the specific choice of γ^p non-essential for this problem. Each distinct near-periodic trajectory within Γ , defined through the above procedure, indicates a distinct vortex-shedding cycle. Moreover, the trajectory exhibited by each $\gamma_j \in \Gamma$ experiences each of the different phases of the vortex-shedding process within a single cycle.

Note that at higher Reynolds numbers, the vortex shedding of flow past a flat plate would lose perfect periodicity due to turbulent fluctuations, though it would retain a strong signature at the characteristic vortex-shedding frequency that is robust to Reynolds number when scaled appropriately (Roshko 1954). For this higher Reynolds number case, one may wish to modify the \mathcal{F} operation to not only compute the lift but also remove the high-frequency components associated with the turbulent fluctuations. This low-frequency signal could be used to establish the (near-)periodic portion of the vortex-shedding behaviour that distinguishes the various cycles.

For the shock problem, the (near-)periodic behaviour is the undulating motion of the shock, with aforementioned finer-scale features such as the supersonic jet exhibiting dynamics about this near-periodic behaviour (cf. figures 1e–h). For this problem, our choice of \mathcal{F} leverages an image filter: an edge-detection algorithm is used to define the upstream (downstream) region of the shock, labelling this region with value zero (unity). An unaltered schlieren image is shown in figure 3(a), with the filtered image using the edge-detection algorithm for that same image shown in figure 3(b). The filtered image is then integrated in space to yield a one-dimensional signal that varies in time, shown in figure 3(c). (The time instance shown in figures 3(a,b) is indicated by the red marker in figure 3(c).) This spatially integrated signal encodes the near-periodic portion of the dynamics, as its value changes as the shock moves within the domain. Similar to the flat plate problem, to indicate phase trajectories undergone by the integrated signal from figure 3(c), figure 3(e) depicts the integrated signal against its time derivative.

Intrinsic phase-based POD

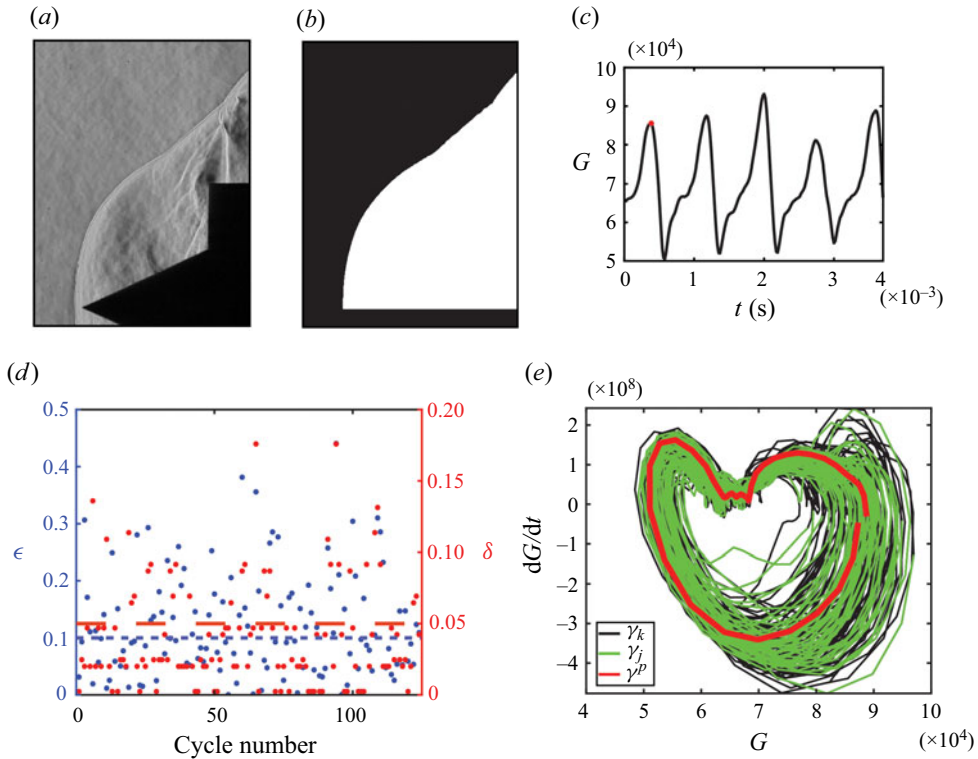


Figure 3. (a) Schlieren image of shock oscillation over cone-cylinder body. (b) Filtered version of (a) to identify phase of near-periodic dynamics. (c) Spatially integrated value of the filtered image shown in (b) plotted against time, with the time instance of (a) marked in red. (d) Nearness of cycle start and end positions (ϵ) and difference of cycle period to nominal period (δ) for all cycles. Threshold values used to determine near-periodic cycles are shown as dashed lines. (e) Cycles plotted using a phase portrait with the value of plot (c) plotted against its time derivative. Cycles that are considered near-periodic using the thresholds of (d) are shown in green, while the representative cycle is shown in red.

Visible in both figures 3(c,e) is an apparent lack of periodicity that nonetheless contains common cyclic behaviour (with significant cycle-to-cycle variation). This cycle-to-cycle variation is demonstrated quantitatively in figure 3(d) through the difference between each cycle's start and end values (ϵ) and the difference between each cycle's period and a reference period (δ). For this problem, the thresholds of these values that are used to define near-periodic orbits were chosen to be $\epsilon^* = 0.1$ and $\delta^* = 0.05$.

These thresholds were chosen as a balance between (1) removing extreme events with significant differences from the nominal trajectory based on closedness and period length, and (2) over-trimming cycles with some differences due to cycle-to-cycle variations that are important to understand. The characteristic periodic cycle γ^p is highlighted in red in figure 3(e), while all cycles are shown in black, and all cycles that are the near-periodic cycles belonging to Γ are shown in green. The computed trajectories Γ contain information about the shock motion for cycles where the undulations remain near a representative periodic orbit. As such, each restricted trajectory $\gamma_j \in \Gamma$ represents a distinct cycle of the shock's oscillation process, and the trajectory exhibited by each of these cycles reflects the various shock locations explored through an oscillation cycle.

Note that the data that we utilize for the shock problem come from schlieren imaging, so that the filtering operation defined above is performed not on the flow state q directly,

but on integrated observations of refractive index changes induced by density gradients in the flow. We do not focus here on this distinction; one could envision an analogous \mathcal{F} operation that acts on the full compressible flow state, identifying discontinuities in any of the flow quantities, and using this location to label upstream and downstream regions and integrate the result over the spatial domain as indicated above.

We highlight that while in both examples the operation \mathcal{F} is chosen to map from the flow state to a real scalar value, this choice is not unique (nor is the choice of mapping, even when restricting the image of the operation to the collection of real numbers). Indeed, mappings onto higher-dimensional outputs might facilitate a more robust analysis of the near-periodic behaviour and the activity driven by this behaviour, described below. As the main focus of this paper is on utilizing a relevant representation of the intrinsic near-periodic dynamics on which to perform POD, we leave the pursuit of an optimal way to extract the underlying near-periodic content as a point for future research.

3.2. A POD formulated for near-periodic driving behaviour

This subsection describes how to use the collection of near-periodic trajectories Γ , associated with the restricted dynamics, to define an expectation operator and inner product that compare distinct cycles along matched phases of the underlying near-periodic content. The result is an energy-ranked modal decomposition, optimal in the usual sense of POD of the best rank- k approximation, for any $k \in \mathbb{Z}^+$, for a data set organized in the proposed phase-matched way.

The expectation operator is defined as an ensemble average over the collection of full-flow-state trajectories that, under the map \mathcal{F} , belong to the near-periodic trajectories Γ . That is, the sample space being drawn from is $j \in N$, and we may, for clarity, explicitly write the full flow state as a random variable or a function of this sample space via

$$\tilde{\mathbf{q}}(t, j) = \mathbf{q}(t), \quad \text{for } t \in [t_{0j}, t_{fj}], j \in N, \tag{3.3}$$

where the dependence on the sample space is stated explicitly.

Intuitively, we seek an expectation operator defined approximately as an ensemble average over these realizations $\langle \tilde{\mathbf{q}} \rangle(t) = (1/n_p) \sum_{j \in N} \tilde{\mathbf{q}}(t, j)$. (Recall that $n_p = \dim(N)$.) One issue with this form is that it is not possible to evaluate this expectation operation at a given time instance, since each $\tilde{\mathbf{q}}(\cdot, j)$ is defined only over the time interval $t \in [t_{0j}, t_{fj}]$. More substantively, even barring that challenge, comparing across the same time within a cycle does not generally correspond to the same phase of the near-periodic dynamics, because of the deviations from periodicity of the nominally periodic behaviour.

To address both of these issues, we define a prototypical periodic trajectory, and relate the time interval of each of the cases with near-periodic cycles $j \in N$ to the nearest time instance of the reference prototypical cycle. This characteristic trajectory is chosen as γ^p , defined in § 3.1, for which the associated $\sqrt{\delta^2 + \epsilon^2}$ is minimized (i.e. minimizing the difference from the nominal period length and the departure from trajectory closedness).

To enable a sensible averaging operation expressed in terms of a unified time interval, we require a map from each time instance $t \in [t_{0j}, t_{fj}]$ to the time instance along the nominal trajectory, $t^p \in [t_0^p, t_f^p]$, for which the near-periodic dynamics are most similar. Denoting $\tau_j = [t_{0j}, t_{fj}]$ and $\tau^p = [t_0^p, t_f^p]$, we define this map $\mathcal{T}_j^p : \tau_j \rightarrow \tau^p$ using

$$t^p = \min_{t' \in [t_0^p, t_f^p]} \|\mathbf{q}_l^p(t') - \mathbf{q}_l(t)\|, \quad t \in \tau_j. \tag{3.4}$$

Intrinsic phase-based POD

There is a related map $\mathcal{T}_p^j : \tau^p \rightarrow \tau_j$, where $\mathcal{T}_j^p = (\mathcal{T}_p^j)^{-1}$ provided that \mathcal{T}_p^j is injective. Intuitively, this requirement is modest: the trajectories are selected to represent nearly the same phenomenon, so supposing a one-to-one mapping from one time interval to the other is natural.

Using (3.4), we arrive at an expression for the expectation operator, given by

$$\langle \tilde{q} \rangle(t) = \frac{1}{n_p} \sum_{j \in N} \tilde{q}(\mathcal{T}_j^p(t), j), \quad t \in \tau^p. \quad (3.5)$$

This expression mirrors a phase average, with the phase of the near-periodic dynamics identified using (3.4). Usually, when phase averages are performed, phase information is known *a priori* through some external forcing such as the rotation of an aerofoil or a rotor. In this case, the phase is identified from the analysis of the instantaneous snapshots: it is this ‘intrinsic phase’ identification and definition that inspired the name IPHaB POD. We will refer to the outcome of (3.5) as the IPHaB mean, to distinguish it from the time mean and from more common phase averages that are based on an externally driven phase variation. (The terminology also fits the name of the method, as the expression (3.5) is formally the mean under the expectation operator used in our proposed IPHaB method.) This definition of an expectation operator is also strongly connected to the conditional expectation defined in Schmidt & Schmid (2019). Indeed, the expression (3.5) entails a conditional expectation operator. Unlike in Schmidt & Schmid (2019), however, the expectation here is conditioned not on a specific event, but rather on a phase of an intrinsic near-periodic behaviour. This difference has required new formulations to the definition of the realizations developed in (3.2) as well as a time mapping that facilitates a clear comparison of the dynamics of each realization relative to the underlying near-periodic behaviour. This time mapping directly affects the definition of the expectation operator in (3.5), and will also manifest itself in the definition of the inner product (described next).

With the expectation operator defined, POD can be performed once a suitable inner product is chosen. Similar to the expectation operator, we emphasize that an inner product that simply integrates along the trajectory of each realization, $\tilde{q}(\cdot, j)$ ($j \in N$), would not account for the imperfect periodicity of the underlying dynamics, and would yield a modal representation that erroneously uses ‘higher modes’ to convey this underlying cyclical behaviour. We therefore choose to perform this trajectory integration relative to a nominal periodic trajectory. Each instance of the trajectory encoded in $\tilde{q}(\cdot, j)$ ($j \in N$) will be cast in terms of the nearest instance on the nominal trajectory. This approach facilitates a meaningful comparison across the various ensembles: all trajectory integrations are aligned with a shared (near-)orbit reflective of the periodic dynamics that each of the ensembles is close to, as defined by (3.2). Denoting the manifold on which the dynamics q evolves as \mathcal{M} , we define the inner product as

$$(r, m) = \int_{t_0^p}^{t_f^p} m^*(\mathcal{T}_j^p(t')) \mathbf{W} r(\mathcal{T}_j^p(t')) dt', \quad \forall r, m \in \mathcal{M}, \quad (3.6)$$

where $(\cdot)^*$ denotes complex conjugation, and \mathbf{W} is an appropriately chosen positive-definite weighting matrix that incorporates quadrature weights to approximate spatial integration of the space-discrete state variables, and allows for the induced norm to be a meaningful representation of physical energy. (Because of the space–time formulation employed here, most applications of this method would likely involve real-valued functions. We leave the complex conjugate in the inner product definition for generality.)

The exact construction of \mathbf{W} varies from problem to problem, but is generally dependent on the field value chosen. The state variable used for the flat plate problem is vorticity, so the weighting chosen is a diagonal matrix with $W_{jj} = \Delta x_j \Delta y_j$, the space between the j th spatial point and its adjacent one. That is, the resulting norm represents the enstrophy (per unit of density) of the system. The ‘state variable’ used in the shock problem is the pixel intensity provided by each schlieren image. The image intensity is driven by density gradients in the flow, and the same weight $W_{jj} = \Delta x_j \Delta y_j$ is chosen. (Note that the third spatial dimension is neglected for both the two-dimensional flat plate problem and for the three-dimensional shock problem, since for the latter, only planewise data are available.) This choice avoids weighting certain regions of space more than others. While the resulting norm has a non-obvious direct physical interpretation because the pixel intensity is not a flow state variable, it is proportional to the square density gradient field.

The eigenvalue problem that yields the IPhaB POD modes then takes the standard form, obtained by searching for the mode ϕ that maximizes the expectation of the projection of the flow state onto that mode. The result is

$$(\mathcal{R}\phi)(t) = \langle (\phi, \tilde{q}') \tilde{q}' \rangle(t) = \int_{t_0}^{t_f} \mathbf{R}(t, t') \mathbf{W} \phi(t') dt' = \lambda \phi(t), \quad t \in \tau^p, \quad (3.7)$$

where, as is common practice, the eigenvalue problem is defined with respect to the fluctuation state variable $\tilde{q}' = \tilde{q} - \langle \tilde{q} \rangle$. Note also that $\mathbf{R}(t, t') = \langle \tilde{q}' \tilde{q}'^* \rangle(t, t')$ is the autocorrelation function.

Because the proposed formulation utilizes the standard POD formulation for a specific choice of inner product and expectation operator, prior results of operator self-adjointness, optimality and mode coherence apply. The operator \mathcal{R} is self-adjoint, and the eigenvalue problem (3.7) yields a collection of modes $\{\phi_1, \phi_2, \dots\}$ that are orthogonal with respect to the inner product (3.6). The associated eigenvalues are positive and real-valued, and are taken here to be ordered such that $\lambda_1 \geq \lambda_2 \geq \dots$ (see e.g. Holmes *et al.* (2012) for a review of these results). An equation for, say, the $n + 1$ most energetic modes can be obtained by maximizing the expectation of the projection of the flow state onto a subspace orthogonal to the span of the first n modes. Optimality of low rank representation and space–time coherence of the modes also hold. For a discussion of the latter point about coherence, and the breakdown of this property when inner products that do not incorporate time are used, see Towne, Schmidt & Colonius (2018). Note that in our case, coherence is with respect to time scaled by the nominal near-periodic trajectory.

3.3. Computing IPhaB POD modes with sampled data

To approximate the time-continuous formulation above with discrete samples of data, we construct a data matrix where the j th column contains the flow state associated with the quasi-periodic trajectory γ_j , $j \in N$. We recall that to facilitate a meaningful comparison across the other trajectories, the eigenvalue problem (3.7) is cast in time according to the nominal periodic trajectory γ^p rather than the j th trajectory. Denoting by n_n the number of time samples of this nominal periodic trajectory, and recalling that $n_p = \dim(N)$ is the number of near-periodic cycles contained within the set Γ , the data matrix to be formed

is then

$$\mathbf{X} = \begin{bmatrix} \tilde{\mathbf{q}}'(\mathcal{T}_{n_1}^p(t_1), n_1) & \tilde{\mathbf{q}}'(\mathcal{T}_{n_2}^p(t_1), n_2) & \dots & \tilde{\mathbf{q}}'(\mathcal{T}_{n_p}^p(t_1), n_p) \\ \tilde{\mathbf{q}}'(\mathcal{T}_{n_1}^p(t_2), n_1) & \tilde{\mathbf{q}}'(\mathcal{T}_{n_2}^p(t_2), n_2) & \dots & \tilde{\mathbf{q}}'(\mathcal{T}_{n_p}^p(t_2), n_p) \\ \vdots & \vdots & \ddots & \vdots \\ \tilde{\mathbf{q}}'(\mathcal{T}_{n_1}^p(t_{n_n}), n_1) & \tilde{\mathbf{q}}'(\mathcal{T}_{n_2}^p(t_{n_n}), n_2) & \dots & \tilde{\mathbf{q}}'(\mathcal{T}_{n_p}^p(t_{n_n}), n_p) \end{bmatrix}. \quad (3.8)$$

Note that each block entry in the matrix \mathbf{X} is defined with respect to the fluctuating quantity. For example, for cycle $n_j \in N$ and time instance $t_k \in \tau_p$, the rows are built as $\tilde{\mathbf{q}}'(\mathcal{T}_{n_j}^p(t_k), n_j) = \tilde{\mathbf{q}}(\mathcal{T}_{n_j}^p(t_k), n_j) - \langle \tilde{\mathbf{q}} \rangle(\mathcal{T}_{n_j}^p(t_k))$.

Using the definition (3.8) of the data matrix, the eigenvalue problem that approximates (3.7) is

$$\mathbf{X}\mathbf{X}^* \tilde{\mathbf{W}}\boldsymbol{\phi}_d = \boldsymbol{\Lambda}\boldsymbol{\phi}_d, \quad (3.9)$$

where $\tilde{\mathbf{W}}$ is a weighting matrix that incorporates the matrix \mathbf{W} from (3.7) as well as quadrature weights to approximate the time integral in the inner product definition. This weighting matrix is block diagonal, and constructed as a concatenation of $(1/n_p)\mathbf{W}$, repeated along block diagonal entries n_p times. This scaling accounts for the ensemble average (and the spatial integral via \mathbf{W}). (For reference, for results shown in later sections for the commonly used space-only POD, which forms the data matrix to have each snapshot of the full data set in a separate column, the weighting matrix is $\tilde{\mathbf{W}} = (1/n_t)\mathbf{W}$, where n_t is the number of snapshots used for the space-only decomposition. This number need not equal $n_n \times n_p$, since generally space-only POD does not omit snapshots from different cycles, and instead uses the full data set.) The eigenvector $\boldsymbol{\phi}_d$ discretely approximates the time-continuous eigenvector $\boldsymbol{\phi}$. With the state vector $\mathbf{q} \in \mathbb{R}^{n_s}$, where n_s is the number of points used to represent the spatial domain Ω , $\boldsymbol{\phi}_d \in \mathbb{R}^{n_s \times n_n}$. The k th collection ($k = 1, \dots, n_n$) of n_s entries in $\boldsymbol{\phi}_d$ approximates $\boldsymbol{\phi}(t_k)$. In practice, the eigenvalue problem (3.9) is typically computationally unwieldy, and the method of snapshots (Sirovich 1987) may be used to extract the desired modes from the smaller eigenvalue problem associated with $\mathbf{X}^* \tilde{\mathbf{W}}\mathbf{X}$.

We now detail the process used to compute $\mathcal{T}_j^p(t_k)$, where $j = n_1, \dots, n_p$, $k = 1, \dots, n_n$. First, we define from the data the nominal near-periodic trajectory $\gamma^p = \{\tilde{\mathbf{q}}(t_{0_l}, l), \tilde{\mathbf{q}}(t_{1_l}, l), \dots, \tilde{\mathbf{q}}(t_{f_l}, l)\}$ for the index $l \in N$ that yields the minimal value of $\sqrt{\epsilon^2 + \delta^2}$. (Note that this trajectory γ^p is the time-sampled variant of the trajectory defined in § 3.2; we avoid using a distinct variable for notational ease.) That is, the time instances indicated in the data matrix (3.8) are defined as $\{t_0, t_1, \dots, t_{n_n}\} := \{t_{0_l}, t_{1_l}, \dots, t_{f_l}\}$. We then build the data matrix where each column contains a different realization of $j \in N, j \neq l$.

For each j th column, we identify $\mathcal{T}_j^p(t_k)$ at each time instance t_k along the nominal near-periodic cycle as follows. First, we identify for the j th cycle the time t_q that yields a restricted flow state nearest to the restricted flow state associated with the nominal periodic cycle:

$$t_q = \arg \max_{t' \in \tau_j} \frac{\mathcal{F}(\tilde{\mathbf{q}})(t_k, l)^* \mathcal{F}(\tilde{\mathbf{q}})(t', j)}{\|\mathcal{F}(\tilde{\mathbf{q}})(t_k, l)\|^2}. \quad (3.10)$$

For some physical problems, including the cone-cylinder shock problem studied in this work, the approach (3.10) can struggle when the near-periodic trajectory of the restricted dynamics contains multiple points in phase space that are close together or overlapping within a single cycle. In the case of the shock problem, this issue occurs when the shock passes the same location, with a similar overall shape, during the upstream and

downstream phases of the shock motion, respectively. In this work, this challenge was managed by further restricting the set of times $t' \in \tau^j$ considered for the optimizer t_q in (3.10) to be within a given range from the prior identified cycle time t_{q-1} . This additional restriction prevents erroneous matching of phases at opposite sides of the dynamic cycle. Future work will aim to resolve this challenge more robustly through a dynamical systems perspective.

Because of finite sampling in time, this nearest flow state in the j th cycle to that associated with the desired phase of the near-periodic dynamics will not in general be desirably close to $\mathcal{F}(\tilde{q})(t_k, l)$. As such, a locally quadratic representation in time, about t_q , is used to find the $t \in [t_{q-1}, t_{q+1}]$ with the maximal projection onto $\mathcal{F}(\tilde{q})(t_k, l)$. Once this time instance $\mathcal{T}_j^p(t_k)$ is determined, quadratic interpolation is used to populate $\tilde{q}'(\mathcal{T}_j^p(t_k), j)$ from the values of the fluctuation flow state $\tilde{q}'(t_{q-1})$, $\tilde{q}'(t_q)$ and $\tilde{q}'(t_{q+1})$. This process is used to construct \mathbf{X} in (3.8), for $j = n_1, \dots, n_p, k = 1, \dots, n_n$.

This proposed approach yields, at each cycle (column in the data matrix \mathbf{X}), snapshots that are matched according to the phase of an underlying near-periodic behaviour driving the dynamics. This approach is useful for problems where the driving behaviour is near-periodic, as it enables a meaningful comparison across cycles whose strong similarities (a common near-periodic behaviour) are obscured by detailed differences. The proposed approach also has utility in cases where the dynamics is perfectly periodic. Even in this setting, finite sampling in time will generally yield a collection of snapshots within a cycle that are at different instances along the phase of the periodic trajectory, which if compared directly do not enable a meaningful assessment of how the dynamics evolve relative to this periodic behaviour.

4. Results

4.1. Vortex shedding from a flat plate with $Re = 100$, $\alpha = 30^\circ$

We first demonstrate IPHaB POD using the $Re = 100$ flat plate problem, described in § 2, with data from Towne *et al.* (2023). We aim to illustrate that when phase-locked on the near-periodic dynamics, the IPHaB mean captures the evolution of the nearly periodic dynamics, and the IPHaB modes contain the evolution of the activity driven by the nearly periodic dynamics. Since this problem has only one periodic feature (vortex shedding), the dynamics is captured in the IPHaB mean, and the IPHaB modes contain only errors due to interpolation. The same data set used to produce figures 1(a–d), and described in § 2, is used for the analysis of this problem.

As discussed in § 3.1, the near-periodic trajectories Γ for the analysis were computed using information about the plate coefficient of lift C_l . Cycles for the lift dynamics were obtained using the local maxima in C_l . The nominal period length of each cycle was computed to be 4.2, and 42 distinct phases were computed along each cycle, i.e. $n_n = 42$. Nine cycles were collected after the initial transient, and due to the system periodicity, all cycles were kept for the IPHaB POD analysis (see § 3.1). In the ensuing discussion, IPHaB POD results are discussed. The results for snapshot POD, SPOD and CSPOD for this data set are shown in Appendix A.

The IPHaB mean and the first two modes of IPHaB are shown in figures 4 and 5, respectively, at different phase instances along the vortex-shedding cycle. The IPHaB mean preserves the leading- and trailing-edge vortex-shedding process that drives the physical system. Four instances of the first and second IPHaB POD modes contain low-intensity activity between the vortices identified in the IPHaB mean. The activity in the IPHaB POD

Intrinsic phase-based POD




Figure 4. Four equally distributed phase instances – (a) 1st, (b) 11th, (c) 22nd, and (d) 33rd, out of a total of 42 distinct phases – of the IPHaB mean for vortex shedding over a flat plate at a stalled angled of attack $\alpha = 30^\circ$. A movie of the flat plate IPHaB mean across all phase instances is available in supplementary movie 3.

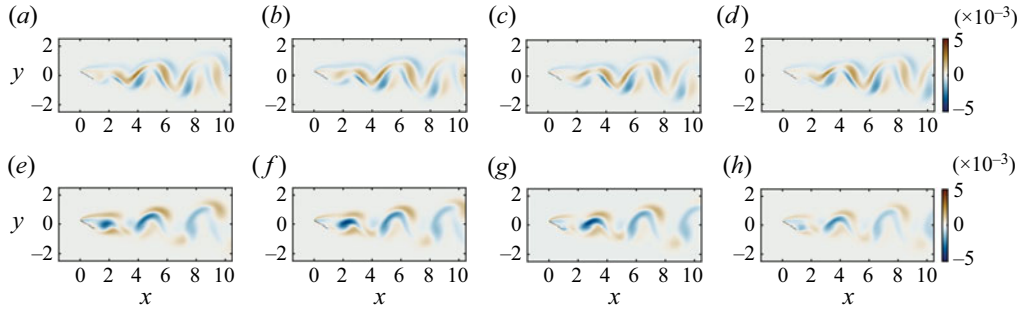


Figure 5. Four equally distributed phase instances – (a,e) 1st, (b,f) 11th, (c,g) 21st, and (d,h) 33rd, out of a total of 44 distinct phases – of (a–d) the first and (e–h) the second IPHaB POD modes for vortex shedding over a flat plate at a stalled angle of attack $\alpha = 30^\circ$. A movie of the flat plate IPHaB POD modes 1 and 2 across all phase instances is available in supplementary movie 4.

modes reflects small cycle-to-cycle numerical errors in the simulation and in interpolating to phase-match the data. This system is driven by a single periodic feature that is fully captured in the IPHaB mean (to within numerical error), and thus makes the activity in the IPHaB modes insignificant. This comprehensive representation of the mean is also indicated by the singular values in figure 6: the left-hand vertical axis contains the non-normalized singular values, and the right-hand vertical axis contains the singular values normalized by the norm of the IPHaB mean. The non-normalized singular values decrease between the second and third modes. When normalized using the norm of the IPHaB mean, the modes contain almost no energy, suggesting that these modes do not contain meaningful structures or activity. The IPHaB POD mode is of size $n_s \times n_n$, while the original data set is of size $n_s \times n_t$, where n_n is again the number of samples in one period, and n_t is the total number of data samples. In general, n_n is much smaller than n_t , representing a reduction in the size of the mode versus the original system.

4.2. Cone-cylinder at $M = 6$

We next perform IPHaB POD analysis on the cone-cylinder shock problem depicted in figures 1(e–h), using the data set from Sasidharan & Duvvuri (2021). To characterize the near-periodic shock motion for this problem described in § 3.1, a Canny edge detection method in the MATLAB 2023b Images Processing tool box was used. Thresholds were set for a sparse identification of the shock front. False shock fronts identified by the algorithm were removed using a recursive moving median filter with window sizes from 3 to 40 time instances. With the locations from the sparse identification of the shock front, we used a spline fit to form a continuous representation of the shock front. Due to the




Figure 6. Singular values for IPHaB for vortex shedding over a flat plate at a stalled angle of attack. The left-hand axis shows the singular values, and the right-hand axis shows the singular values normalized by the norm of the ensemble-averaged mean $E := \|\langle \tilde{W}\tilde{q} \rangle\|^2$.




Figure 7. Four equally distributed phase instances – (a) 1st, (b) 12th, (c) 23rd, and (d) 34th, out of 44 phases – of the IPHaB mean are shown. A movie of the cone-cylinder IPHaB POD modes 1–4 across all phase instances is available in supplementary movie 5.

axisymmetric nature of the cone-cylinder body, all spatial positions below the axis of symmetry were labelled zero (evident in the filtered image in figure 3b). We smoothed the spatially integrated signal using a Savitzky Golay filter. We determined the period length and the start and end of each cycle based on local maxima of G . The period length varied from 6.4912×10^{-4} to 8.9474×10^{-4} s, with average period length $\bar{T} = 7.8763 \times 10^{-4}$ s. Based on the thresholds ϵ^* and δ^* stated in § 3.1 for this problem, 48 out of 124 cycles were determined to be near-periodic, with the nominal periodic cycle γ^p satisfying $\delta^p = 0.0199$ and $\epsilon^p = 0.2321$.

The IPHaB mean (cf. figures 7a–d) represents the time-varying motion of the shock front, which consists of a conical shock, a separation shock and a bow shock, over a single representative cycle. The shock front dynamics of the mode reflects the dynamics observed in the data, but is simpler than the original data, with fewer tethered dynamics observed. The shock front often appears as a sharp edge in the mean, but in some time instances it is more spatially distributed. Time instances where the shock is less sharp

(e.g. [figure 7d](#)) reflect that there are significant cycle-to-cycle variations in the spatial location of the shock in that part of the cycle, while time instances where the shock appears very sharp (e.g. [figures 7b,c](#)) reflect that the shock positioning is very consistent across cycles. [Figures 7\(b,c\)](#) are phases of the cycle where the separation bubble grows under the adverse pressure gradient induced by the coalescence of the cone and cylinder shocks. We learn from the sharpness of the mode that this process is highly regular and repeatable across cycles. [Figures 7\(a,d\)](#), which show slightly blurrier shock fronts, correspond to the collapse and reformation of the separation bubble and subsequently of the shock system. We learn from the blurriness of the mode that these processes are less regular and more chaotic cycle to cycle. Note that the more non-periodic cycles are included by increasing the threshold values of δ^* and ϵ^* , the blurrier these regions become, as the cycles become increasingly disparate. Conversely, the overall mode becomes less converged with too few cycles included.

The IPHaB POD modes contain details of tethered physical processes, whose dynamics follow the cycle of the shock front. Tethered dynamics, such as in the supersonic jet and the triple point, is visible in the separation region, particularly in the initial phases of the cycle, including the 1st (cf. [figures 8a,e,i,m](#)) and 12th (cf. [figures 8b,f,j,n](#)) instances along the cycle. The movement of the supersonic jet and the triple point upwards, away from the axis of symmetry can be seen in the first three instances shown of the first ([figures 8a–c](#)), second ([figures 8e–g](#)), third ([figures 8i–k](#)) and fourth ([figures 8m–o](#)) modes. The non-normalized singular values and the singular values normalized based on the norm of the ensemble-averaged mean are shown in [figure 9](#). A small decrease can be seen between the second and third singular values. The normalized singular values shown on the axis on the right show that the IPHaB POD modes contain a small amount of energy, suggesting that the IPHaB mean contains most of the energy from the flow field.

Three more instances of the first IPHaB POD mode zoomed in near the cone-cylinder body are shown in [figure 10](#). The supersonic jet is marked by black ellipses, and the activity near the base cylinder is marked by red ellipses. As the supersonic jet moves away from the axis of symmetry, the strength of the jet also weakens, indicated by the decrease in the activity in the mode. The location of the strongest activity downstream of the shock front also changes at different phases of the shock pulsation. In the first phase instance (cf. [figure 10a](#)), the strongest activity is present near the wall of the base cylinder and from the wall of the cone all the way to the shock front. At the second phase instance (cf. [figure 10b](#)), the activity primarily remains near the wall of the base cylinder, though the activity decreases above the base cylinder shoulder, and the separation bubble associated with the boundary layer along the cone wall starts developing. In the last phase instance (cf. [figure 10c](#)), the separation bubble is near its largest size, and the supersonic jet becomes less coherent. The strong growth in bubble size from [figure 10\(b\)](#) to [figure 10\(c\)](#) demonstrates that the separation process is dominated by the final quarter of the undulation cycle.

While these distinct tethered processes are prominent across modes 1–4, it should also be noted that features of the IPHaB mean are also inherited in modes near the early and late parts of the cycle. These artefacts are evident in, for example, [figure 8\(d\)](#), where large-wavelength flow structures dominate the mode's behaviour. These phase instances where the large-scale shock is evident in the modes correspond with the instances when the cycle-to-cycle variations in the shock motion are most pronounced. These cycle variations require that changes to large-scale shock features be conveyed by the IPHaB modes, in addition to the tethered processes of primary interest. This imperfect isolation of the near-periodic and tethered dynamics is a possible issue to be addressed in future work,

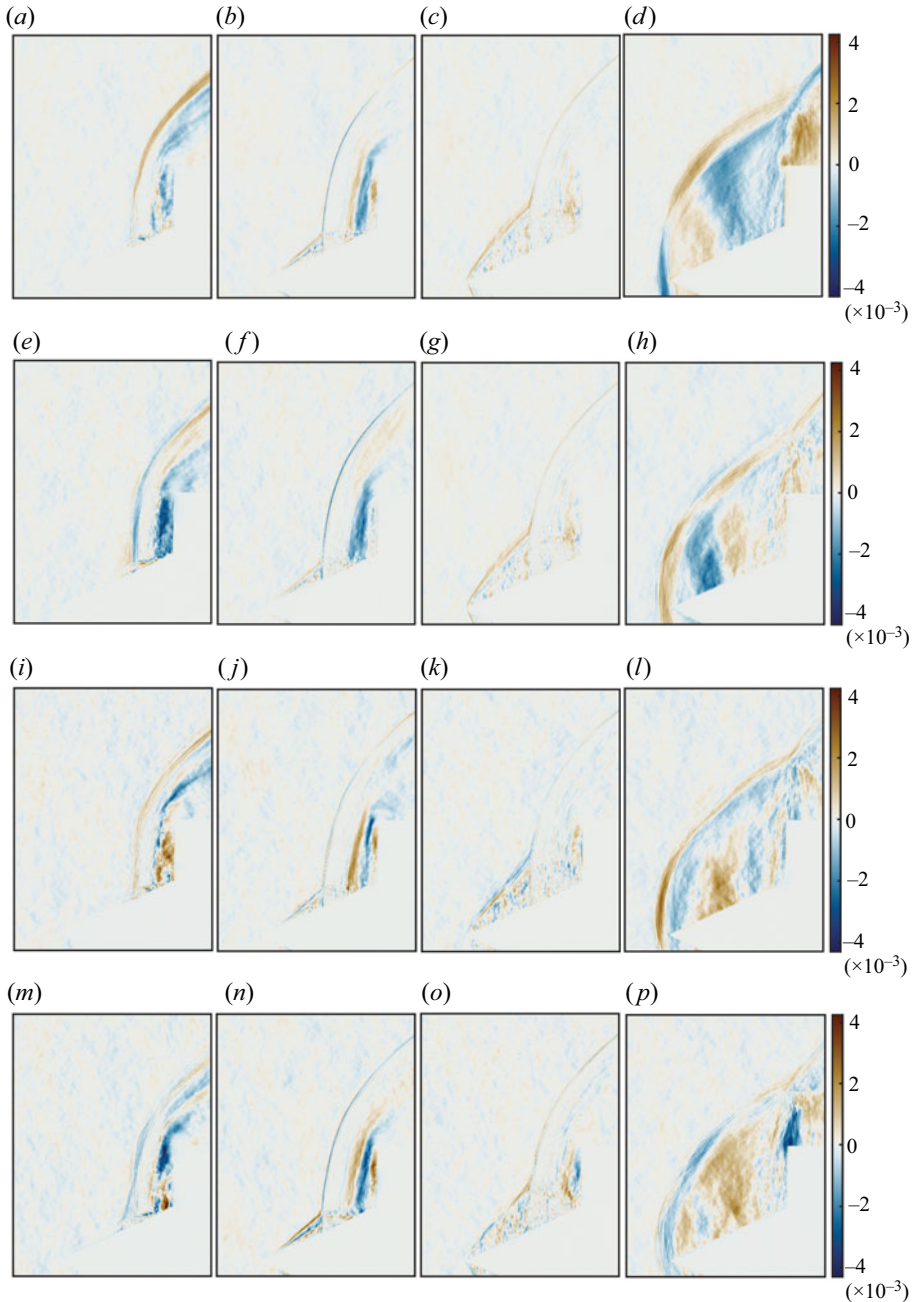


Figure 8. (a–d) The first, (e–h) the second, (i–l) the third, and (m–p) the fourth IPHaB POD modes for shock pulsations over a cone-cylinder body. For each mode, four equally distributed phase instances – (a,e,i,m) 1st, (b,f,j,n) 12th, (c,g,k,o) 23rd, and (d,h,l,p) 34th, out of 44 phases – are shown. A movie of the cone-cylinder IPHaB POD modes 1–4 across all phase instances is available in supplementary movie 6.

though we also highlight that this outcome provides information about instances when cycle-to-cycle variations are more and less pronounced.

The modes are observed to be fairly insensitive to noise. The data contain significant high-frequency noise, which is not observed to dominate the features of the first four

Intrinsic phase-based POD




Figure 9. Analogue to figure 6, for the cone-cylinder shock problem.

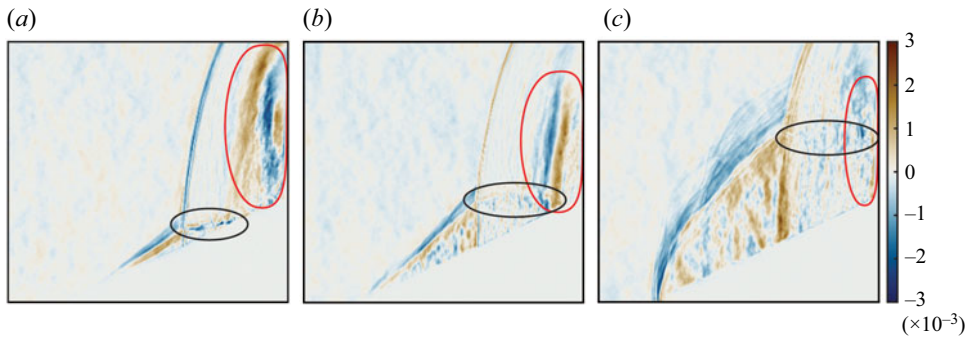


Figure 10. The (a) 7th, (b) 17th, and (c) 27th phase instances of IPHaB POD first mode zoomed in to highlight the evolution of the supersonic jet (black ellipses) and the activity near the base cylinder (red ellipses).

modes in figure 8. Additionally, if the low-frequency noise mentioned in § 2 is not mitigated, then these modes remain consistent but an additional mode is observed that accounts for the low-frequency noise. Future work will consider more systematically the influence of noise on the output of the method.

5. Comparison to other methods

Here, we perform space-only POD, SPOD (using the method of Schmidt 2022) and CSPOD (using the method of Heidt *et al.* 2023) on the pulsating shock data. In this section, we provide a concise comparison of only the mean and first mode representations at a single phase and/or frequency, as is relevant for the given method. Additional results for each method are provided in Appendix A for reference, along with details of the implementation of these methods for these data.

The IPHaB mean and first IPHaB mode are shown at a single phase in figures 11(a,b), respectively. The time mean is provided in figure 11(c), along with the first mode of




Figure 11. A comparison of the mean and first mode of IPhaB POD, space-only POD, SPOD and CSPOD for the shock pulsation problem. (a) The IPhaB mean and (b) the first mode are shown at one phase instance. (c) The time mean. The first mode of (d) space-only POD is shown, along with (e) the real part of the first mode at the strongest frequency for SPOD, and (f) one phase of the first mode of CSPOD.

space-only POD (figure 11d). The real part of the first mode at the strongest frequency is shown for SPOD (figure 11e), and a single phase of the first mode of CSPOD (figure 11f) for the shock data. More modes are provided in Appendix A, and movies of the SPOD and CSPOD modes are provided in supplementary movies 9 and 10. In figure 11, we observe that the IPhaB mean (figure 11a) and first mode (figure 11b) are representing distinct dynamics at the chosen phase. The mean (figure 11a) represents the location of the shock system at that one phase, including a sharp bow shock associated with the cylinder and a sharp oblique shock associated with the cone. The mode (figure 11b) shows turbulent dynamics near the cone surface, in the supersonic jet, and larger structures near the cylinder wall that occur at this phase of the dynamic cycle. The time mean in figure 11(c) shows fewer details than the phase-localized IPhaB mean, such as a lack of the oblique shock near the cone. The POD, SPOD and CSPOD modes (figures 11d–f) show features of the global shock behaviour at multiple phases of its motion: the shock signature can be observed at the tip of the cone as well as at the midpoint of the cone simultaneously. These modes are not representative of the shock’s instantaneous behaviour, but rather represent its global dynamics by spatially mixing multiple phases. Multiple such modes must be added with various weightings to reconstruct a single instance in the dynamic cycle. We expect that for a perfectly periodic system, CSPOD would not artificially mix phases of the shock dynamics into a single phase instance of its mode, but for imperfectly periodic systems, CSPOD is observed to generate modes that mix multiple phases of the global dynamics into a single phase of the mode. None of the methods shows the

small-scale turbulent fluctuations of the supersonic jet seen in the first mode of IPhaB POD in [figure 11\(b\)](#).

To augment the simple comparison highlighted in [figure 11](#), several space-only POD, SPOD and CSPOD modes are provided for both the flat plate and shock problems in [Appendix A](#). We highlight some of the more relevant conclusions from that appendix here. For the flat plate problem, CSPOD has a strong first mode that, when added to the time mean, is approximately equivalent to the IPhaB mean (see [figure 17](#)). The next modes have smaller singular values, similar to the IPhaB modes of small energy relative to the mean (see [figure 18](#)). Both POD and SPOD require many modes to represent the vortex-shedding cycle from the flat plate, with a slow drop off in singular values (see [figures 13–16](#)). For the shock problem, none of the modes for POD, SPOD or CSPOD, including at higher frequencies for SPOD, is observed to efficiently represent the tethered dynamics of the IPhaB POD first mode. A continuous rather than precipitous drop off in singular values is observed for all other methods applied to the shock problem, while the IPhaB singular values are observed to be many orders of magnitude smaller than the IPhaB mean.

The POD and SPOD are not able to represent the tethered turbulent features because neither has a mechanism to represent dynamics that is local-in-phase within a global cycle. The POD and SPOD modes are global modes that are driven to optimally represent the system at all times, leading them to generate modal structures that combine behaviour over multiple phases of the flow cycle. The CSPOD is able to represent tethered features for perfectly periodic systems, and gives an equivalent outcome to IPhaB POD for the flat plate wake problem. But for problems with imperfect periodicity, CSPOD mixes distinct phases that are not matched in the near-periodic dynamics, resulting in artefacts from the near-cyclical behaviour into higher modes.

The singular values are provided for each method in [Appendix A](#). The IPhaB method aims to contain as much of its global dynamics as possible in its mean, leaving the modes for fine-scale and local-in-phase dynamics. Therefore we look for low-magnitude singular values compared to our mean. Conversely, space-only POD and SPOD aim to represent global dynamics in the modes, and should therefore be expected to have larger singular values for their modes in relation to the norm of the mean. The singular values of the CSPOD first mode should be much larger than its other singular values, as its first mode is most similar to the IPhaB mean. For imperfectly periodic systems, the drop off between the first mode and higher modes would be expected to narrow as the near-periodic dynamics is increasingly represented in higher modes.

6. Conclusion

A proper orthogonal decomposition (POD) framework was developed, IPhaB POD, that is appropriate for data with periodic or near-periodic cycles. For these systems, the proposed method leverages a dynamical systems representation of a set of restricted dynamics sampled from the full dynamical state. These restricted dynamics convey information about the global dynamic cycle of the flow. For example, for the flat plate problem considered in this paper, the restricted dynamics were of the lift coefficient, which distinguished the phases of the vortex-shedding process. We provided a framework that uses these restricted dynamics to identify which cycles of the full system dynamics are associated with a characteristic near-periodic orbit (as opposed to those corresponding to cases where the dynamics undergo an extreme trajectory). All cycles that contained a characteristic near-periodic trajectory were collected into a set Γ and were considered for

the POD analysis; others were discarded. We also used this framework to define the cycle from the data that contains a most typical near-periodic trajectory, γ^p .

An averaging operation and inner product were then defined that referred each instance from any one of the trajectories in Γ to the nearest corresponding instance in γ^p . This process allowed for the various cycles within Γ , which each contained meaningfully similar near-periodic dynamics but with cycle-to-cycle differences, to be compared. Once the inner product and expectation were defined that allowed an appropriate comparison of the various near-periodic trajectories in Γ , an IPhaB mean was defined, which is a phase average that uses a comparison between different cycles of the data to define an implicit phase. With this IPhaB mean subtracted, a typical POD analysis was then performed using the standard eigenvalue problem of a suitably constructed data matrix.

The method was tested on two test problems. The first problem was from a high-fidelity computation of a canonical $Re = 100$ flow past an aerofoil at the large angle of attack $\alpha = 30^\circ$ (Towne *et al.* 2023). In this laminar flow, the only behaviour is of periodic vortex shedding that alternates from the leading and trailing edges. The IPhaB POD captured this vortex shedding within the IPhaB mean. The IPhaB POD modes represented interpolation errors. This simple problem highlights that IPhaB POD captures global near-periodic or periodic dynamics in the IPhaB mean, with the modes reserved for other tethered dynamics if they exist. No other dynamics exist in the flat plate problem, so no additional modes were needed to physically represent the problem. The singular values of the modes, compared to a norm of the IPhaB mean, were shown to be very small, highlighting their relative unimportance.

The second problem was an experimental data set from schlieren data of high Reynolds number and Mach number flow past a cone-cylinder system (Sasidharan & Duvvuri 2021). The IPhaB POD represented the global, near-periodic shock dynamics within the IPhaB mean, and yielded clearly identifiable tethered dynamics in the IPhaB POD modes. These higher modes represented the boundary layer separation as well as the motion of the triple point and its associated supersonic jet that are known to be key processes in this physical system (Sasidharan & Duvvuri 2021). Moreover, these higher modes evolved with the shock motion, providing more physically intuitive representations of instantaneous flow features for this complex system.

The outcome of applying IPhaB POD to the shock problem was briefly compared to the outcome of applying space-only POD, spectral POD and cyclostationary spectral POD. The other three methods were found to produce modes that represented multiple phases of the shock motion at a single instance of a mode. No mode observed from the other three methods was found to efficiently represent the tethered dynamics, including the supersonic jet, and its phase relationship with the near-periodic shock dynamics.

Supplementary movies. Supplementary movies are available at <https://doi.org/10.1017/jfm.2024.1061>.

Acknowledgements. We would like to acknowledge Professor S. Duvvuri at the Indian Institute of Science, and his graduate student V. Sasidharan, for providing us with the cone-cylinder schlieren data used for analysis.

Funding. This work was funded by the National Science Foundation under grant 2118209.

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

Author ORCIDs.

 Akhileshwar Borra <https://orcid.org/0000-0002-2932-9403>;

 Andres Goza <https://orcid.org/0000-0002-9372-7713>;

 Theresa Saxton-Fox <https://orcid.org/0000-0003-1328-4148>.

Intrinsic phase-based POD




Figure 12. Time average of vortex shedding over a flat plate at stall angle of attack.




Figure 13. (a) First, (b) second, (c) third, and (d) fourth space-only POD modes for vortex shedding over a flat plate at stall angle of attack.

Appendix A. Space-only POD, SPOD and CSPOD

Here, we provide the results for the space-only POD, SPOD and CSPOD methods in greater detail. These results are provided for both the flat plate (§ A.1) and cone-cylinder shock problems (§ A.2). The details for the data sets for each problem are given in § 2. For each problem, we provide the mean quantities, which are subtracted from the data matrix associated with each method, as well as some modes for each method. A few results that are most salient for comparison with the proposed IPHaB POD method are provided in § 5.

A.1. Flat plate

The time-averaged mean for the flat plate data is shown in figure 12. This mean is subtracted from the data prior to calculating the space-only POD, SPOD and CSPOD modes. The first four modes of space-only POD applied to the flat plate data are shown in figure 13. Space-only POD represents the alternating leading- and trailing-edge vortex shedding across a cascade of modes. The first and second modes (cf. figures 13a,b) show regions of vorticity that are approximately up–down symmetric about $y = 0$, with a matched streamwise length scale that is a nearly constant spatial wavelength. These modes are identical apart from a streamwise shift in vorticity to capture convection of the shed vortices. The third mode (cf. figure 13c) and fourth mode (cf. figure 13d) likewise come in pairs to convey convection of the length scales of the vortex-shedding process captured by these two modes. The singular values for the modes show a gradual decrease until approximately 80 modes in figure 14, indicating that many modes are required to reconstruct the vortex-shedding process.

The SPOD was implemented using the algorithm from Schmidt (2022) and was adapted to fit this data set. This algorithm uses Welch’s method to estimate the spectral density of the signal to go from a time domain to a frequency domain. A Hamming window size of 125 snapshots and overlap 63 was used, resulting in 5 blocks and thus 5 modes at 64 frequencies. Similar to space-only POD, the time-averaged mean was subtracted from the




Figure 14. Singular values of space-only POD modes for vortex shedding over a flat plate at stall angle of attack.

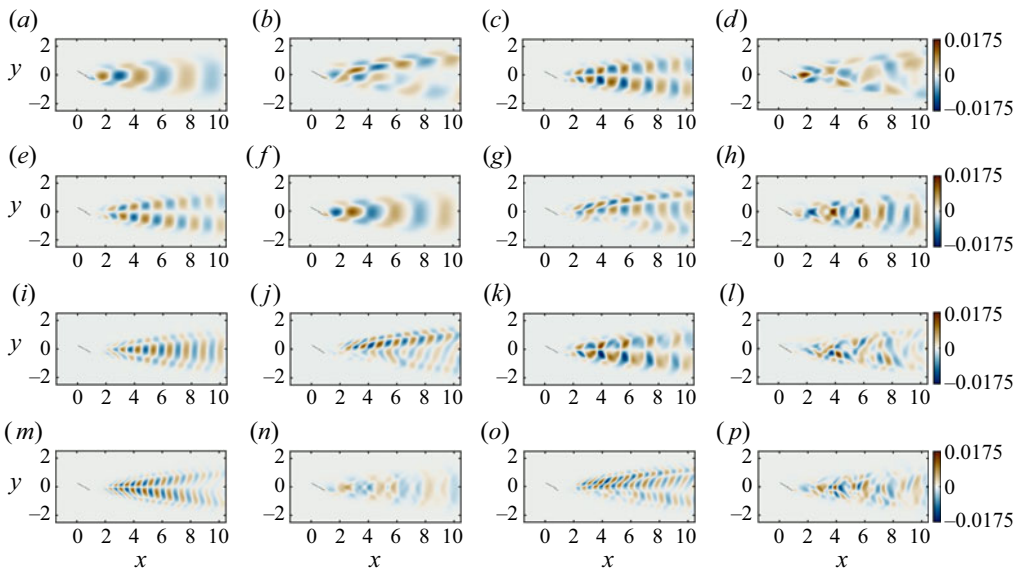


Figure 15. The SPOD modes at (a–d) 0.24 Hz, (e–h) 0.48 Hz, (i–l) 0.72 Hz, and (m–p) 0.96 Hz for vortex shedding over a flat plate at stall angle of attack. The (a,e,i,m) first, (b,f,j,n) second, (c,g,k,o) third, and (d,h,l,p) fourth modes are shown for the specified frequencies. Movies of the first four SPOD modes at the four frequencies at different phases are available in supplementary movie 7.

data set prior to performing SPOD. Figure 15 shows the first four SPOD modes for the four most dominant frequencies of the flat plate data: 0.24, 0.48, 0.72, 0.96 Hz. Many of the SPOD modes look similar to space-only POD modes shown in figure 13. The singular values, shown in figure 16, decay slowly with frequency, indicating that modes at many frequencies are required to represent the vortex shedding.

Intrinsic phase-based POD




Figure 16. Singular values of SPOD for vortex shedding over a flat plate at stall angle of attack.

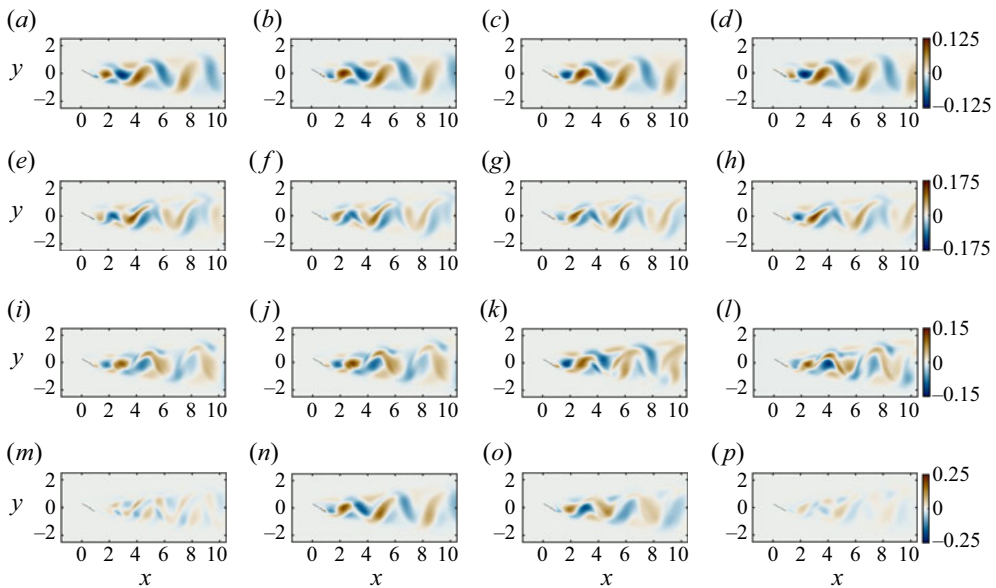


Figure 17. Four equally distributed phase instances – (a,e,i,m) 1st, (b,f,j,n) 11th, (c,g,k,o) 21st, and (d,h,l,p) 33rd, out of a total of 44 distinct phases – of the (a–d) first, (e–h) second, (i–l) third, and (m–p) fourth CSPOD modes for vortex shedding over a flat plate at a stalled angle of attack $\alpha = 30^\circ$. A movie of the first four flat plate CSPOD modes across all phase instances is available in supplementary movie 8.

Figure 17 shows the first four CSPOD modes at four equally spaced time instances within the vortex-shedding cycle. The CSPOD modes were calculated using the methods and algorithms of Heidt & Colonius (2024) and Heidt *et al.* (2023). To manage the discrepancy between the frequency with which data is saved and the vortex-shedding frequency, interpolation is used to ensure that the data are as periodic as possible.




Figure 18. Singular values of CSPOD for vortex shedding over a flat plate at stall angle of attack.




Figure 19. Time average of Mach 6 flow over the cone-cylinder body.

The period \bar{T} is provided to the method, which is equivalent to the nominal cycle length for IPHaB POD. The CSPOD method produces modes at various frequency shift values γ , but here we report only the results for the most energetic frequency, $\gamma = 0$. The transform itself is computed with a Hamming window of length $5\bar{T}$ and overlap 0.75. Like the space-only POD, SPOD and IPHaB results, the mean is subtracted from the data before decomposition. For the CSPOD method, this is the window mean, or the time-averaged mean for each window. (Note that the window means are not shown here but are nearly identical to the time average shown in [figure 12](#).) The CSPOD modes exhibit physical features of instantaneous vortex shedding that evolve in time across each cycle instance. There is a significant drop-off between the first and second singular values, shown in [figure 18](#), indicating that the first mode for CSPOD is dominant. The CSPOD is capable of representing the vortex shedding almost entirely through the sum of its mean and its first mode, similar to how IPHaB POD can represent these data with its IPHaB mean.

Intrinsic phase-based POD




Figure 20. The (a) first, (b) second, (c) third, and (d) fourth space-only POD modes for Mach 6 flow over the cone-cylinder body.




Figure 21. Singular values of space-only POD for Mach 6 flow over the cone-cylinder body.

A.2. Cone-cylinder at $M = 6$

Figures 19 and 20 respectively show the time-averaged mean and first four space-only POD modes for the cone-cylinder problem. The singular values for space-only POD are shown in figure 21. In all four modes, two features are consistently present: the shock front from the time-averaged mean, and the bow shock front. The modes mix these features together, rather than showing them at distinct phase instances, to create a globally optimal solution for summatively recreating every snapshot from the same set of modes. We can also notice a lack of tethered features throughout any of the modes; all of the dominant modes are being summed to recreate the global near-periodic dynamics of the shock.

The SPOD code from Schmidt (2022) was modified for the cone-cylinder data. All the time instances were used when performing SPOD, and the time-averaged mean was subtracted. For Welch's method, a Hamming window of 512 snapshots with overlap 256 snapshots was used and resulted in a total of 20 blocks. Thus the flow field was decomposed into 257 frequencies, each with 20 modes. Figure 22 contains the first four SPOD modes for frequencies 1224, 2561, 3785, 5121 Hz, while figure 23 shows the singular values for all the SPOD modes. Clear peaks are observed in figure 23,

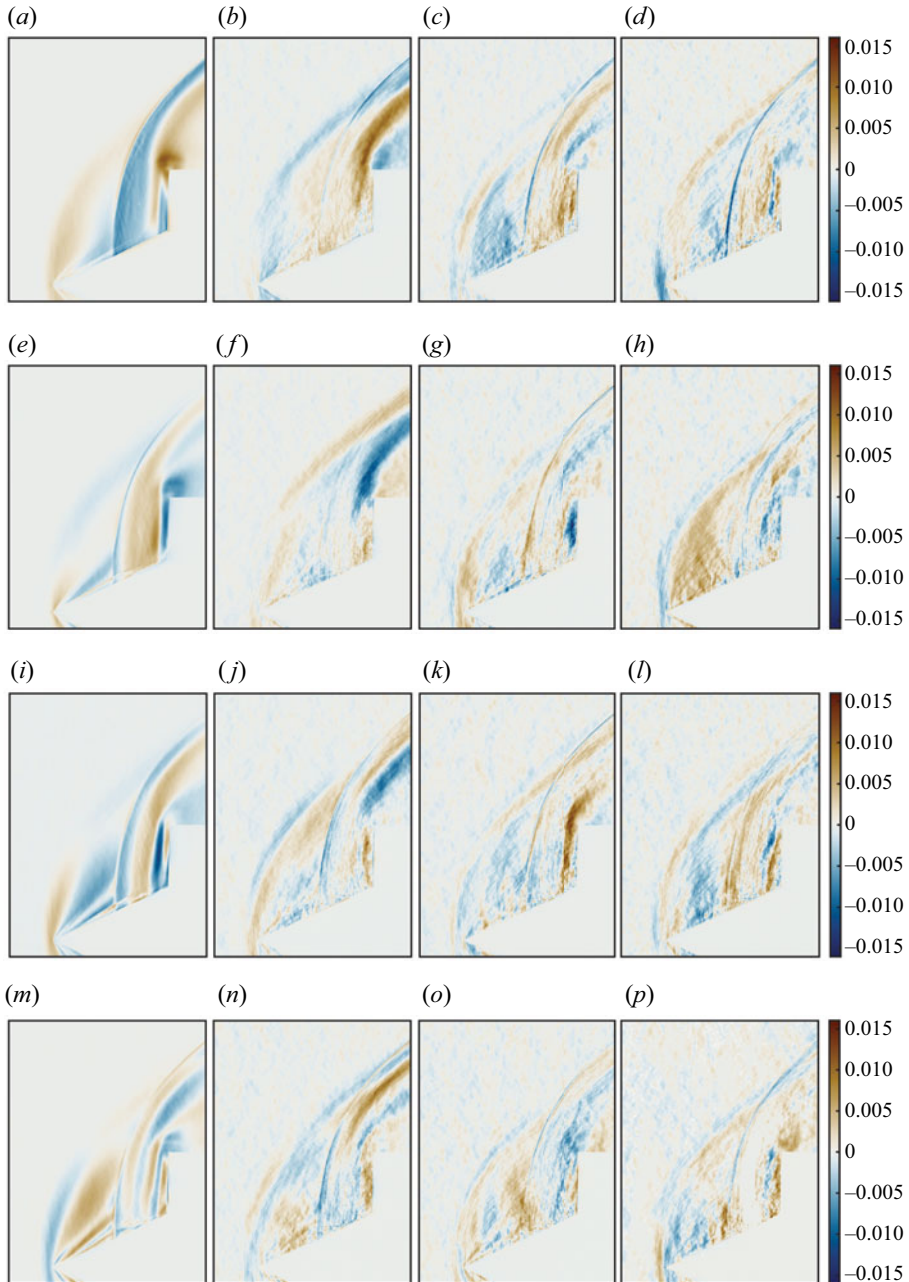


Figure 22. The SPOD modes at (a–d) 1224 Hz, (e–h) 2561 Hz, (i–l) 3785 Hz, and (m–p) 5121 Hz for Mach 6 flow over the cone-cylinder body. The (a,e,i,m) first, (b,f,j,n) second, (c,g,k,o) third and (d,h,l,p) fourth modes are shown for the specified frequencies. Movies of the first four SPOD modes at the four frequencies at different phases are available in supplementary movie 9.

Intrinsic phase-based POD




Figure 23. Singular values of SPOD for Mach 6 flow over the cone-cylinder body.

reflecting the dominant frequency of the shock cycle and its harmonics. The first mode at the dominant frequencies is observed to be two orders of magnitude larger than the higher modes. These first modes consistently show features of the global shock cycle at various phases of its evolution (figures 22a,e,i,m), rather than showing evidence of tethered dynamics such as the supersonic jet. The higher modes at these frequencies show more spatial variability, but do not show clear definition of local-in-phase turbulent dynamics.

Figure 24 shows the first four CSPOD modes for frequency shift $\gamma = 0$ at four equally spaced instances for the dynamic shock data. In contrast to the flat plate problem, the data are not phase-locked due to the more significant variation between cycles, making that phase-locking highly non-trivial (and a key focus of the IPhaB approach). The period for the nominal CSPOD cycle is chosen to be the average period length \bar{T} , but the window length is set to $11\bar{T}$ to ensure accurate spectral estimates while keeping an overlap of 0.75. As in the case of the flat plate problem, the window means for the cone-cylinder are extremely close to the time average in figure 19. This causes the CSPOD modes to exhibit the same time-averaged shock front observed in both space-only POD and SPOD across all modes. The CSPOD modes do show some evidence of phase variation, as observed in figure 24 with the oblique shock on the cone starting very near the cone in figure 24(a), and appearing farther from the cone in figure 24(b). But this phase variation is not synchronized across the shock dynamics: the bow shock does not show an obviously matched phase evolution, multiple positions of the shock are superimposed in each phase of the modes, and the evolution is not consistent with physical observations as one continues to progress into figures 24(c,d). Instead, figures 24(c) and 24(d) appear very similar to figures 24(a) and 24(b), but with flipped sign. None of the modes shows evidence of the tethered, local-in-phase dynamics that was observed in the first mode of IPhaB POD. The singular values for the modes are shown in figure 25. The singular values are all of the same order of magnitude, rather than showing a sharp drop off, as was seen for the flat plate data.

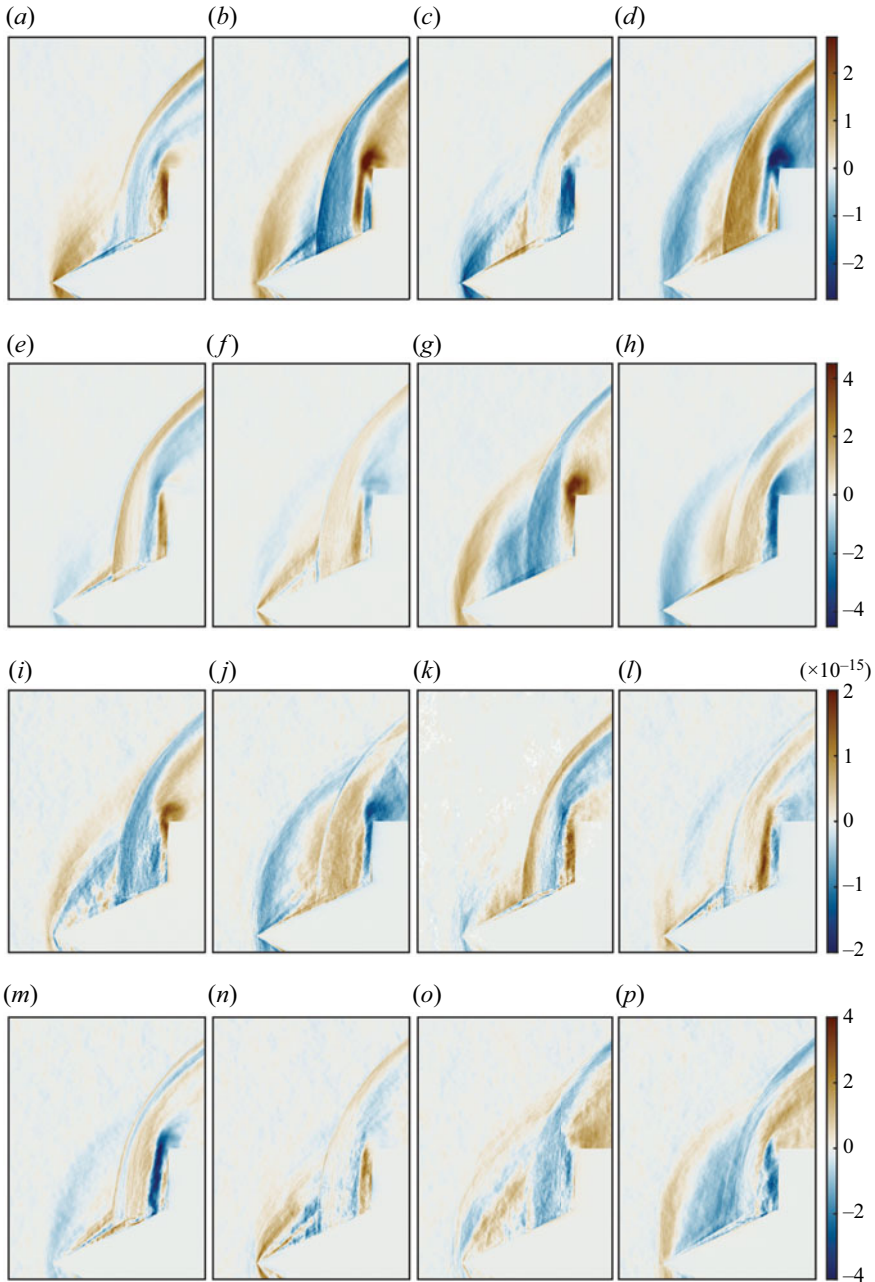


Figure 24. The (a–d) first, (e–h) second, (i–l) third, and (m–p) fourth CSPOD modes for shock pulsations over the cone-cylinder body. For each mode, four equally distributed phase instances – (a,e,i,m) 1st, (b,f,j,n) 12th, (c,g,k,o) 23rd, and (d,h,l,p) 34th, out of 44 phases – are shown. Movies of the first four cone-cylinder IPhaB POD modes across all phase instances are available in supplementary movie 10.

Intrinsic phase-based POD




Figure 25. Singular values of CSPOD for shock pulsations over the cone-cylinder body.

REFERENCES

- AHUJA, S. & ROWLEY, C. 2010 Feedback control of unstable steady states of flow past a flat plate using reduced-order estimators. *J. Fluid Mech.* **645**, 447–478.
- BERGMANN, M., CORDIER, L. & BRANCHER, J.-P. 2005 Optimal rotary control of the cylinder wake using proper orthogonal decomposition reduced-order model. *Phys. Fluids* **17** (9), 097101.
- BORRA, A. & SAXTON-FOX, T. 2022 Proper orthogonal decomposition for a turbulent boundary layer with aero-optic distortion. *AIAA J.* **60** (6), 3555–3564.
- COLONIUS, T. & TAIRA, K. 2008 A fast immersed boundary method using a nullspace approach and multi-domain far-field boundary conditions. *Comput. Meth. Appl. Mech. Engng* **197** (25–28), 2131–2146.
- DESIKAN, S.L.N., PANDIAN, S., SIMHA, P.P. & SAHOO, N. 2022 Experiments on flow past cavity and the application of modal decomposition techniques. *Eur. J. Mech. (B/ Fluids)* **94**, 293–298.
- FENG, L.-H., WANG, J.-J. & PAN, C. 2011 Proper orthogonal decomposition analysis of vortex dynamics of a circular cylinder under synthetic jet control. *Phys. Fluids* **23** (1), 014106.
- FLORYAN, D. & GRAHAM, M.D. 2021 Discovering multiscale and self-similar structure with data-driven wavelets. *Proc. Natl Acad. Sci.* **118** (1), e2021299118.
- FRANCESCHINI, L., SIPP, D., MARQUET, O., MOULIN, J. & DANDOIS, J. 2022 Identification and reconstruction of high-frequency fluctuations evolving on a low-frequency periodic limit cycle: application to turbulent cylinder flow. *J. Fluid Mech.* **942**, A28.
- FREUND, J.B. & COLONIUS, T. 2009 Turbulence and sound-field POD analysis of a turbulent jet. *Intl J. Aeroacoust.* **8** (4), 337–354.
- GORDEYEV, S.V. & THOMAS, F.O. 2013 A temporal proper decomposition (TPOD) for closed-loop flow control. *Exp. Fluids* **54**, 1–16.
- GOZA, A. & COLONIUS, T. 2018 Modal decomposition of fluid–structure interaction with application to flag flapping. *J. Fluids Struct.* **81**, 728–737.
- HEIDT, L. & COLONIUS, T. 2024 Spectral proper orthogonal decomposition of harmonically forced turbulent flows. *J. Fluid Mech.* **985**, A42.
- HEIDT, L., COLONIUS, T., NEKKANTI, A., SCHMIDT, O.T., MAIA, I. & JORDAN, P. 2023 Cyclostationary analysis of forced turbulent jets. In *AIAA AVIATION 2023 Forum*, p. 3652.
- HILBERG, D., LAZIK, W. & FIEDLER, H.E. 1994 The application of classical POD and snapshot POD in a turbulent shear layer with periodic structures. *Appl. Sci. Res.* **53**, 283–290.
- HOLMES, P., LUMLEY, J.L., BERKOOZ, G. & ROWLEY, C.W. 2012 Turbulence, coherent structures, dynamical systems and symmetry. In *Proper Orthogonal Decomposition*, pp. 68–105. Cambridge University Press.
- HOSSAIN, S., BERGSTROM, D.J. & CHEN, X.B. 2015 Visualisation and analysis of large-scale vortex structures in three-dimensional turbulent lid-driven cavity flow. *J. Turbul.* **16** (10), 901–924.
- HUANG, C., ANDERSON, W.E., HARVAZINSKI, M.E. & SANKARAN, V. 2016 Analysis of self-excited combustion instabilities using decomposition techniques. *AIAA J.* **54** (9), 2791–2807.

- KOURENTIS, L. & KONSTANTINIDIS, E. 2012 Uncovering large-scale coherent structures in natural and forced turbulent wakes by combining PIV, POD, and FTLE. *Exp. Fluids* **52**, 749–763.
- KUTZ, J.N., FU, X. & BRUNTON, S.L. 2016 Multi-resolution dynamic mode decomposition. *SIAM J. Appl. Dyn. Syst.* **15** (2), 713–735.
- LENGANI, D., SIMONI, D., ZUNINO, P. & BERTINI, F. 2017 Experimental investigation on the time–space evolution of a laminar separation bubble by proper orthogonal decomposition and dynamic mode decomposition. *J. Turbomach.* **139** (3), 031006.
- MENDEZ, M.A., BALABANE, M. & BUCHLIN, J.-M. 2019 Multi-scale proper orthogonal decomposition of complex fluid flows. *J. Fluid Mech.* **870**, 988–1036.
- MURRAY, N., SÄLLSTRÖM, E. & UKEILEY, L. 2009 Properties of subsonic open cavity flow fields. *Phys. Fluids* **21** (9), 095103.
- PADOVAN, A., OTTO, S.E. & ROWLEY, C.W. 2020 Analysis of amplification mechanisms and cross-frequency interactions in nonlinear flows via the harmonic resolvent. *J. Fluid Mech.* **900**, A14.
- PADOVAN, A. & ROWLEY, C.W. 2022 Analysis of the dynamics of subharmonic flow structures via the harmonic resolvent: application to vortex pairing in an axisymmetric jet. *Phys. Rev. Fluids* **7** (7), 073903.
- PICKERING, E., RIGAS, G., NOGUEIRA, P.A.S., CAVALIERI, A.V.G., SCHMIDT, O.T. & COLONIUS, T. 2020 Lift-up, Kelvin–Helmholtz and Orr mechanisms in turbulent jets. *J. Fluid Mech.* **896**, A2.
- PING, H., ZHU, H., ZHANG, K., ZHOU, D., BAO, Y., XU, Y. & HAN, Z. 2021 Dynamic mode decomposition based analysis of flow past a transversely oscillating cylinder. *Phys. Fluids* **33** (3), 033604.
- RAVINDRAN, S. 2000 Reduced-order adaptive controllers for fluid flows using POD. *J. Sci. Comput.* **15**, 457–478.
- RIBERIO, J. & WOLF, W. 2017 Identification of coherent structures in the flow past a NACA0012 airfoil via proper orthogonal decomposition. *Phys. Fluids* **29** (8), 085104.
- ROSHKO, A. 1954 On the drag and shedding frequency of two-dimensional bluff bodies. *Tech. Rep.*. NACA Technical Note 3169.
- ROWLEY, C. & DAWSON, S. 2017 Model reduction of flow analysis and control. *Annu. Rev. Fluid Mech.* **49**, 387–417.
- SASIDHARAN, V. & DUVVURI, S. 2021 Large- and small-amplitude shock-wave oscillations over axisymmetric bodies in high-speed flow. *Journal of Fluid Mechanics Rapids* **913**, R7.
- SAXTON-FOX, T., LOZANO-DURÁN, A. & MCKEON, B.J. 2022 Amplitude and wall-normal distance variation of small scales in turbulent boundary layers. *Phys. Rev. Fluids* **7** (1), 014606.
- SAXTON-FOX, T. & MCKEON, B.J. 2017 Coherent structures, uniform momentum zones and the streamwise energy spectrum in wall-bounded turbulent flows. *J. Fluid Mech.* **826**, R6.
- 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.
- SCHMIDT, O. 2022 Spectral proper orthogonal decomposition using multitaper estimates. *Theor. Comput. Fluid Dyn.* **36**, 741–754.
- SCHMIDT, O.T. & SCHMID, P.J. 2019 A conditional space–time POD formalism for intermittent and rare events: example of acoustic bursts in turbulent jets. *J. Fluid Mech.* **867**, R2.
- SCHMIDT, O.T., TOWNE, A., RIGAS, G., COLONIUS, T. & BR’ES, G.A. 2018 Spectral analysis of jet turbulence. *J. Fluid Mech.* **855**, 953–982.
- SEENA, A. & SUNG, H.J. 2011 Dynamic mode decomposition of turbulent cavity flows for self-sustained oscillations. *Intl J. Heat Fluid Flow* **32**, 1098–1110.
- SIROVICH, L. 1987 Turbulence and the dynamics of coherent structures part I: coherent structures. *Q. Appl. Math.* **45**, 561–571.
- SZUBERT, D., GROSSI, F., JIMINEZ GARCIA, A., HOARAU, Y., HUNT, J. & BRAZA, M. 2015 Shock–vortex shear-layer interaction in the transonic flow around a supercritical airfoil at high Reynolds number in buffet conditions. *J. Fluids Struct.* **55**, 276–302.
- TAIRA, K., BRUNTON, S.L., DAWSON, S.T.M., ROWLEY, C.W., COLONIUS, T., MCKEON, B.J., SCHMIDT, O.T., GORDEYEV, S., THEOFILIS, V. & UKEILEY, L.S. 2017 Modal analysis of fluid flows: an overview. *AIAA J.* **55** (12), 4013–4041.
- TAIRA, K. & COLONIUS, T. 2007 The immersed boundary method: a project method. *J. Comput. Phys.* **255** (2), 2118–2137.
- TAYLOR, J. & GLAUSER, M. 2004 Towards practical flow sensing and control via POD and LSE based low-dimensional tools. *J. Fluids Engng* **126** (3), 337–345.
- TOWNE, A., *et al.* 2023 A database for reduced-complexity modeling of fluid flows. *AIAA J.* **61** (7), 2867–2892.

Intrinsic phase-based POD

- TOWNE, A., SCHMIDT, O.T. & COLONIUS, T. 2018 Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis. *J. Fluid Mech.* **847**, 821–867.
- WEIGHTMAN, J., AMILI, O., HONNERY, D., SORIA, J. & EDGINTON-MITCHELL, D. 2018 Signatures of shear-layer unsteadiness in proper orthogonal decomposition. *Exp. Fluids* **59**, 180.
- WU, Y., TANG, Z., YANG, S., SKOTE, M., TANG, H., ZHANG, G. & SHAN, Y. 2017 Proper-orthogonal-decomposition study of turbulent near wake of S805 airfoil in deep stall. *AIAA J.* **55** (6), 1959–1969.
- ZHANG, Q., LIE, Y. & WANG, S. 2014 The identification of coherent structures using proper orthogonal decomposition and dynamic mode decomposition. *J. Fluids Struct.* **49**, 53–72.
- ZHANG, Y., HODŽIĆ, A., EVRARD, F., VAN WACHEM, B. & VELTE, C.M. 2023 Phase proper orthogonal decomposition of non-stationary turbulent flow. *Phys. Fluids* **35** (4), 045109.