

Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis

Aaron Towne^{1,†}, Oliver T. Schmidt² and Tim Colonius²

¹Center for Turbulence Research, Stanford University, Stanford, CA 94305, USA

²California Institute of Technology, Pasadena, CA 91125, USA

(Received 3 August 2017; revised 13 March 2018; accepted 30 March 2018)

We consider the frequency domain form of proper orthogonal decomposition (POD), called spectral proper orthogonal decomposition (SPOD). Spectral POD is derived from a space–time POD problem for statistically stationary flows and leads to modes that each oscillate at a single frequency. This form of POD goes back to the original work of Lumley (*Stochastic Tools in Turbulence*, Academic Press, 1970), but has been overshadowed by a space-only form of POD since the 1990s. We clarify the relationship between these two forms of POD and show that SPOD modes represent structures that evolve coherently in space and time, while space-only POD modes in general do not. We also establish a relationship between SPOD and dynamic mode decomposition (DMD); we show that SPOD modes are in fact optimally averaged DMD modes obtained from an ensemble DMD problem for stationary flows. Accordingly, SPOD modes represent structures that are dynamic in the same sense as DMD modes but also optimally account for the statistical variability of turbulent flows. Finally, we establish a connection between SPOD and resolvent analysis. The key observation is that the resolvent-mode expansion coefficients must be regarded as statistical quantities to ensure convergent approximations of the flow statistics. When the expansion coefficients are uncorrelated, we show that SPOD and resolvent modes are identical. Our theoretical results and the overall utility of SPOD are demonstrated using two example problems: the complex Ginzburg–Landau equation and a turbulent jet.

Key words: computational methods, low-dimensional models, turbulent flows

1. Introduction

Coherent flow features, or structures, play an important role in turbulent flows. This has led to efforts to extract these structures from data as well as to model them using simplified equations. Typically, this involves defining a set of modes that compactly describe the structures. Some of the most widely used techniques are summarized in recent reviews by Rowley & Dawson (2017) and Taira *et al.* (2017).

[†] Email address for correspondence: atowne@stanford.edu

Fifty years after its introduction by Lumley (1967, 1970), proper orthogonal decomposition (POD) remains one of the most widely used techniques for educing coherent structures from flow data. The method is known in other disciplines by a variety of names, including principal component analysis and Karhunen–Loève decomposition. Proper orthogonal decomposition seeks a set of deterministic modes that optimally capture the energy, or variance, of an ensemble of stochastic flow data. A rigorous statement of this objective involves defining the ensemble of interest in terms of flow data and choosing an associated inner product and expectation operator; these choices determine the properties of the modes that are obtained.

This paper focuses on a specific form of POD called spectral proper orthogonal decomposition (SPOD). To be clear, we are not referring to the method recently proposed by Sieber, Paschereit & Oberleithner (2016) that was given the same name. Instead, we are using the terminology introduced by Picard & Delville (2000) to describe a space–time formulation of POD for statistically stationary flows that goes back to Lumley (1967, 1970). This terminology is motivated by the fact that the method involves decomposition of the cross-spectral density tensor and leads to modes that each oscillate at a single frequency. Spectral POD has been applied to a variety of flows, including pipes (Hellström & Smits 2014), boundary layers (Tutkun & George 2017), mixing layers (Delville *et al.* 1999; Braud *et al.* 2004), jets (Glauser, Leib & George 1987; Arndt, Long & Glauser 1997; Citriniti & George 2000; Gordeyev & Thomas 2000; Gudmundsson & Colonius 2011; Schmidt *et al.* 2017a), wakes (Tutkun, Johansson & George 2008; Araya, Colonius & Dabiri 2017) and the flow around an airfoil (Abreu, Cavalieri & Wolf 2017).

Since its introduction and popularization by Sirovich (1987), Aubry *et al.* (1988) and Aubry (1991), a different form of POD has come to dominate the literature. This form of POD decomposes the spatial correlation tensor and leads to spatially orthogonal modes that are modulated in time by expansion coefficients with random time dependence. In what follows, we refer to this form of POD as space-only POD.

Several factors appear to have contributed to the dominance of space-only POD. First, SPOD requires time-resolved data that were not available using particle image velocimetry until recently. On the other hand, it is laborious to obtain the cross-spectral densities for a large number of spatial points using hot wires, and large arrays of hot wires (Citriniti & George 2000) become intrusive. In terms of simulation data, SPOD requires relatively long-time integrations that would have been prohibitive using computational resources available at that time. Finally, space-only POD provides an appealing basis for Galerkin projection of the Navier–Stokes equations, which became popular with the rise of the dynamical systems perspective on turbulence (e.g. Aubry *et al.* 1988; Holmes *et al.* 1997; Noack *et al.* 2003).

Perhaps due to the dominance of the space-only variant, there is a lack of clarity in the literature regarding the relationship between space-only and spectral POD. A survey of some of the most cited and/or recent review articles that address POD reveals a wide range of perspectives. Some articles describe POD in purely abstract terms and make no distinction between the two versions of the method (Berkooz, Holmes & Lumley 1993; Liang *et al.* 2002). Many papers exclusively present space-only POD (Moin & Moser 1989; Sirovich 1989; Holmes *et al.* 1997; Chatterjee 2000; Rowley, Colonius & Murray 2004; Cordier & Bergmann 2008; Rowley & Dawson 2017), while one early review (George 1988) considers only spectral POD. There are some articles that treat the two variants as separate methods, but the relationship between them is not explored in detail (Aubry *et al.* 1988; Picard & Delville 2000; Chen & Kareem 2005; Tinney & Jordan 2008; Holmes *et al.* 2012; Taira *et al.* 2017).

Finally, one recent review describes space-only POD as an approximation of spectral POD (George 2017).

Compounding this lack of clarity is the inconsistent and incompatible use of the terms ‘classical POD’ and ‘snapshot POD’. Some authors use these terms to refer to the methods we are calling spectral and space-only POD respectively (e.g. Hellström & Smits 2014; Mula & Tinney 2014; George 2017). Others use the name ‘classical POD’ to refer to space-only POD and ‘snapshot POD’ to refer to a particular computational shortcut for computing space-only POD modes (e.g. Hilberg, Lazik & Fiedler 1994; Pinier *et al.* 2007; Cordier & Bergmann 2008).

The first part of this paper seeks to clarify the relationship between the space-only and spectral formulations of POD. We will show that they are fundamentally different from one another – whereas it is often stated that both methods identify coherent structures, only SPOD modes evolve coherently in space and time. This suggests that SPOD is better suited for identifying physically meaningful coherent structures in stationary flows. We also derive formulae relating space-only and spectral POD eigenvalues and eigenvectors.

The second part of the paper establishes a connection between SPOD and dynamic mode decomposition (DMD). Dynamic mode decomposition was developed by Schmid (2010) as an alternative to POD for identifying coherent structures from flow data with the specific aim of obtaining modes that describe the flow dynamics, i.e. the evolution of the flow from one time instant to the next. This objective was motivated in part by criticism of space-only POD, specifically that the averaging process used to obtain the spatial correlation tensor causes important dynamical information about the flow to be lost. Our analysis affirms that this criticism is well founded for space-only POD but shows that it does not apply to SPOD. Moreover, we will show that SPOD modes are in fact optimally averaged DMD modes obtained from an ensemble DMD problem for stationary flows.

Several other methods have been proposed in recent years to try to bridge the gap between the spatial orthogonalization of space-only POD and the temporal orthogonalization of DMD. Cammilleri *et al.* (2013) proposed Cronos–Koopman analysis by treating the projection coefficients from space-only POD as observables in a Koopman analysis (approximated by DMD). The aforementioned method of Sieber *et al.* (2016) filters the temporal correlation tensor over a time horizon, leading to an *ad hoc* interpolation between space-only POD, which is obtained when the filter width is zero, and the discrete Fourier transform, which is recovered when the filter width is the entire interval of the data. Noack *et al.* (2016) developed a method called recursive DMD (RDMD) which combines features of space-only POD and DMD. The first RDMD mode is given by the DMD mode that minimizes the time-averaged residual between the modal expansion and the data. Subsequent modes achieve the same objective under the constraint of orthogonality with previous modes. This leads to modes that each oscillate at a single frequency and are spatially orthogonal to all other modes at all frequencies. The unique properties of each of these methods make them useful for different purposes. While a detailed comparison is beyond the scope of this paper, our analysis shows that SPOD is optimal by construction for the task of identifying flow structures that evolve coherently in both space and time.

The third part of the paper shows that SPOD is closely related to resolvent analysis. Resolvent analysis (also called input/output analysis and frequency response analysis) has its roots in linear systems and control theory. The resolvent operator is derived from linearized flow equations and constitutes a transfer function between inputs and outputs of interest. It has been used to study the linear response of flows to

external body forces and perturbations (Trefethen *et al.* 1993; Farrell & Ioannou 2001; Schmid & Henningson 2001; Jovanović & Bamieh 2005; Bagheri *et al.* 2009; Sipp *et al.* 2010) and to forcing from the nonlinear terms in the Navier–Stokes equations (McKeon & Sharma 2010; Sharma & McKeon 2013). In the latter context, the method can be derived by partitioning of the Navier–Stokes equations into terms that are linear and nonlinear with respect to perturbations to the turbulent mean flow. Resolvent analysis then identifies frequency-dependent modes that are optimal in terms of their linear gain between the nonlinear terms and the output. The idea is to then use a small set of the highest-gain modes as a basis for the output.

The connection we draw between SPOD and resolvent analysis is based on a new statistical interpretation of the resolvent-mode reconstruction of turbulent flows. Due to the sensitivity of the Navier–Stokes equations to small perturbations, each realization of a turbulent flow, e.g. a different run of the same experiment, produces a unique time history that cannot be reliably predicted by knowledge of the time history of a different realization. In other words, turbulent flows are random in the sense defined by Landahl & Mollo Christensen (1992) and Pope (2000), among others. Accordingly, a statistical description that accounts for many such trajectories provides more information about the likely properties of any specific realization. This leads to a statistical interpretation of the expansion coefficients that are used in the resolvent-mode reconstruction of the flow, which is a departure from past studies that have described them as deterministic quantities described entirely by their amplitude and phase.

We will show that SPOD and resolvent modes are identical when these expansion coefficients are uncorrelated, which is typically associated with white-noise forcing. This can be viewed as a statistical counterpart to the relationship between DMD and resolvent analysis recently shown by Sharma, Mezić & McKeon (2016). More generally, we will demonstrate the importance of properly accounting for the cross-correlations between the expansion coefficients. We will show that if the expansion coefficients are not treated as statistical quantities, the optimal reconstruction of the flow is always governed by the leading SPOD mode at each frequency; thus, the quality of the approximation is dependent first and foremost on the low-rank nature of the cross-spectral density tensor rather than the resolvent operator. This limitation can be overcome by using SPOD modes to estimate the statistics of the expansion coefficients.

The remainder of the paper is organized as follows. Section 2 describes and compares the space-only and spectral formulations of POD. Section 3 outlines a procedure for estimating SPOD modes using time-resolved flow data. The relationship between SPOD and DMD is explored in §4. Section 5 establishes a connection between SPOD and resolvent analysis, and shows the key role of the statistics of the resolvent-mode expansion coefficients. Two example problems that demonstrate the relationships between the various decompositions are given in §6, and §7 summarizes and concludes the paper.

2. Proper orthogonal decomposition

The basic objective underlying POD is this: given a zero-mean stochastic process $\{q(\mathbf{z}; \xi)\}$, find the deterministic function $\phi(\mathbf{z})$ that best approximates the stochastic function on average (Lumley 1967, 1970). Here, \mathbf{z} is a set of independent variables and ξ is an element in the probability space that parameterizes the stochastic variable. We assume that each realization of the stochastic process belongs to a Hilbert space

\mathcal{H} with inner product $\langle \cdot, \cdot \rangle$ and define $E\{\cdot\}$ to be the expectation operator over the probability space. With these definitions, this objective is formalized by maximizing the quantity

$$\lambda = \frac{E\{|\langle \mathbf{q}(\mathbf{z}; \xi), \boldsymbol{\phi}(\mathbf{z}) \rangle|^2\}}{\langle \boldsymbol{\phi}(\mathbf{z}), \boldsymbol{\phi}(\mathbf{z}) \rangle} \quad (2.1)$$

over all $\boldsymbol{\phi}(\mathbf{z}) \in \mathcal{H}$. That is, we wish to find the deterministic function $\boldsymbol{\phi}(\mathbf{z})$ that maximizes the expected value of the normalized projection of the stochastic function.

A standard variational approach can be used to show that the function $\boldsymbol{\phi}(\mathbf{z})$ that maximizes (2.1) must satisfy the eigenvalue problem

$$\langle \mathbf{C}(\mathbf{z}, \mathbf{z}'), \boldsymbol{\phi}(\mathbf{z}') \rangle^* = \lambda \boldsymbol{\phi}(\mathbf{z}), \quad (2.2)$$

where

$$\mathbf{C}(\mathbf{z}, \mathbf{z}') = E\{\mathbf{q}(\mathbf{z}; \xi) \mathbf{q}^*(\mathbf{z}'; \xi)\} \quad (2.3)$$

is the two-point correlation tensor. Throughout this paper, we use an asterisk superscript to denote both the complex conjugate of a scalar and the Hermitian transpose of a vector or tensor. The properties of the solutions of the eigenvalue problem (2.2) depend critically on the properties of the kernel \mathbf{C} , which in turn depend on the definition of the stochastic ensemble. Fluid flows are described by space–time fields $\mathbf{q}(\mathbf{x}, t)$, and the space-only and spectral variants of POD are obtained by using these flow data to define the stochastic ensemble and the associated inner product and averaging operation in two different ways, as described in the following sections.

2.1. Space-only POD

The most commonly employed form of POD generates spatial modes $\boldsymbol{\phi}(\mathbf{x})$. This is accomplished by defining the stochastic ensemble to consist of snapshots of the flow field at different time instances. In other words, the flow at each instant is treated as a realization of a stochastic process. The appropriate inner product is then

$$\langle \mathbf{u}, \mathbf{v} \rangle_x = \int_{\Omega} \mathbf{v}^*(\mathbf{x}, t) \mathbf{W}(\mathbf{x}) \mathbf{u}(\mathbf{x}, t) d\mathbf{x}, \quad (2.4)$$

where \mathbf{u} and \mathbf{v} are any two elements in \mathcal{H} , Ω denotes the spatial domain over which the flow is defined, and the weight \mathbf{W} is a positive-definite Hermitian tensor of appropriate dimension. We will restrict our attention to bounded spatial domains, but note that unbounded homogeneous dimensions can be accommodated by transforming those directions to Fourier space (Lumley 1967, 1970; George 2017). The expectation operator for this definition of the stochastic ensemble is simply a time average, so we are restricted to statistically stationary flows.

The quantity to maximize is

$$\lambda = \frac{E\{|\langle \mathbf{q}(\mathbf{x}, t), \boldsymbol{\phi}(\mathbf{x}) \rangle_x|^2\}}{\langle \boldsymbol{\phi}(\mathbf{x}), \boldsymbol{\phi}(\mathbf{x}) \rangle_x} \quad (2.5)$$

and the resulting Fredholm eigenvalue problem is

$$\int_{\Omega} \mathbf{C}(\mathbf{x}, \mathbf{x}') \mathbf{W}(\mathbf{x}') \boldsymbol{\phi}(\mathbf{x}') d\mathbf{x}' = \lambda \boldsymbol{\phi}(\mathbf{x}), \quad (2.6)$$

where

$$\mathbf{C}(\mathbf{x}, \mathbf{x}') = E\{\mathbf{q}(\mathbf{x}, t)\mathbf{q}^*(\mathbf{x}', t)\} \quad (2.7)$$

is the two-point spatial correlation tensor. This tensor is a nuclear kernel, i.e. it is compact and $\int_{\Omega} \mathbf{C}(\mathbf{x}, \mathbf{x}) d\mathbf{x} < \infty$. As a result, Hilbert–Schmidt theory guarantees that the eigenmodes satisfying (2.6) have a number of special properties. First, there exists a countably infinite set of eigenmodes, $\{\phi_j, \lambda_j\}$, that can be ranked according to their eigenvalue, $\lambda_1 \geq \lambda_2 \geq \dots \geq 0$. Each eigenvalue gives the average energy captured by that mode, measured in the spatial norm induced by the inner product (2.4), and the total energy of the flow is given by the sum of the eigenvalues.

The eigenvectors are orthogonal, $\langle \phi_j, \phi_k \rangle_x = \delta_{jk}$, and provide a complete basis for \mathbf{q} . Accordingly, the flow field can be expanded as

$$\mathbf{q}(\mathbf{x}, t) = \sum_{j=1}^{\infty} a_j(t) \phi_j(\mathbf{x}), \quad (2.8)$$

with $a_j(t) = \langle \mathbf{q}(\mathbf{x}, t), \phi_j(\mathbf{x}) \rangle_x$. This expansion is optimal in its ability to capture the flow energy; if the expansion is truncated at order n , any other orthogonal expansion of the same order will capture less energy. The expansion coefficients are uncorrelated at zero time lag,

$$E\{a_j(t)a_k^*(t)\} = \lambda_j \delta_{jk}. \quad (2.9)$$

Finally, the eigenmodes provide a diagonal representation of the two-point spatial correlation tensor

$$\mathbf{C}(\mathbf{x}, \mathbf{x}') = \sum_{j=1}^{\infty} \lambda_j \phi_j(\mathbf{x}) \phi_j^*(\mathbf{x}') \quad (2.10)$$

and are therefore its principal components. Accordingly, space-only POD modes optimally represent spatial correlations within the flow.

2.2. Spectral proper orthogonal decomposition

Alternatively, we can seek modes that depend on both space and time. This is accomplished by defining the stochastic ensemble to consist of a collection of realizations of the time-dependent flow. For example, different runs of the same experiment are considered to be realizations of a stochastic process. The appropriate inner product is then

$$\langle \mathbf{u}, \mathbf{v} \rangle_{x,t} = \int_{-\infty}^{\infty} \int_{\Omega} \mathbf{v}^*(\mathbf{x}, t) \mathbf{W}(\mathbf{x}) \mathbf{u}(\mathbf{x}, t) d\mathbf{x} dt \quad (2.11)$$

and the expectation operator is an ensemble average over different stochastic realizations of the flow. The quantity to maximize is then

$$\lambda = \frac{E\{|\langle \mathbf{q}(\mathbf{x}, t), \phi(\mathbf{x}, t) \rangle_{x,t}|^2\}}{\langle \phi(\mathbf{x}, t), \phi(\mathbf{x}, t) \rangle_{x,t}}, \quad (2.12)$$

which leads to the eigenvalue problem

$$\int_{-\infty}^{\infty} \int_{\Omega} \mathbf{C}(\mathbf{x}, \mathbf{x}', t, t') \mathbf{W}(\mathbf{x}') \phi(\mathbf{x}', t') d\mathbf{x}' dt' = \lambda \phi(\mathbf{x}, t), \quad (2.13)$$

where

$$\mathbf{C}(\mathbf{x}, \mathbf{x}', t, t') = E\{\mathbf{q}(\mathbf{x}, t)\mathbf{q}^*(\mathbf{x}', t')\} \quad (2.14)$$

is the two-point space–time correlation tensor.

Because they persist indefinitely, statistically stationary flows have infinite energy in a space–time norm. As a result, the space–time correlation tensor (2.14) is not nuclear and the eigenmodes of (2.13) do not possess the properties generally associated with POD. To remedy this, a new eigenvalue problem can be obtained in spectral space from which modes with useful properties can be obtained. In what follows, we focus exclusively on stationary flows.

To derive the spectral eigenvalue problem, we recall that the correlation tensor of a wide-sense stationary flow depends only on the difference between two times,

$$\mathbf{C}(\mathbf{x}, \mathbf{x}', t, t') \rightarrow \mathbf{C}(\mathbf{x}, \mathbf{x}', t - t'). \quad (2.15)$$

Then, the cross-spectral density tensor \mathbf{S} can be defined as the Fourier transform pair of the correlation tensor,

$$\mathbf{S}(\mathbf{x}, \mathbf{x}', f) = \int_{-\infty}^{\infty} \mathbf{C}(\mathbf{x}, \mathbf{x}', \tau) e^{-i2\pi f\tau} d\tau \quad (2.16)$$

and

$$\mathbf{C}(\mathbf{x}, \mathbf{x}', \tau) = \int_{-\infty}^{\infty} \mathbf{S}(\mathbf{x}, \mathbf{x}', f) e^{i2\pi f\tau} df. \quad (2.17)$$

Using these definitions, the following result can be derived: for any frequency f' , the function $\boldsymbol{\phi}(\mathbf{x}, t) = \boldsymbol{\psi}(\mathbf{x}, f') e^{i2\pi f' t}$ is a solution of the eigenvalue problem (2.13) with eigenvalue $\lambda(f')$, where $\boldsymbol{\psi}(\mathbf{x}, f')$ and $\lambda(f')$ satisfy the spectral eigenvalue problem

$$\int_{\Omega} \mathbf{S}(\mathbf{x}, \mathbf{x}', f') \mathbf{W}(\mathbf{x}') \boldsymbol{\psi}(\mathbf{x}', f') d\mathbf{x}' = \lambda(f') \boldsymbol{\psi}(\mathbf{x}, f'). \quad (2.18)$$

This result was first given by Lumley (1967, 1970); we offer an alternative derivation in appendix A.

The cross-spectral density tensor is nuclear, so the eigenmodes of (2.18) at each frequency inherit properties analogous to those of space-only POD modes. There is a countably infinite set of eigenfunctions $\boldsymbol{\psi}_j(\mathbf{x}, f)$ at each frequency that are orthogonal to all other modes at the same frequency in the spatial inner product (2.4), i.e. $\langle \boldsymbol{\psi}_j, \boldsymbol{\psi}_k \rangle_x = \delta_{jk}$. The Fourier modes of each flow realization are optimally expanded as

$$\hat{\mathbf{q}}(\mathbf{x}, f) = \sum_{j=1}^{\infty} a_j(f) \boldsymbol{\psi}_j(\mathbf{x}, f), \quad (2.19)$$

with $a_j(f) = \langle \hat{\mathbf{q}}(\mathbf{x}, f), \boldsymbol{\psi}_j(\mathbf{x}, f) \rangle_x$, and the expansion coefficients are uncorrelated,

$$E\{a_j(f) a_k^*(f)\} = \lambda_j(f) \delta_{jk}. \quad (2.20)$$

The cross-spectral density tensor has the diagonal representation

$$\mathbf{S}(\mathbf{x}, \mathbf{x}', f) = \sum_{j=1}^{\infty} \lambda_j(f) \boldsymbol{\psi}_j(\mathbf{x}, f) \boldsymbol{\psi}_j^*(\mathbf{x}', f), \quad (2.21)$$

so the SPOD modes are its principal components. Furthermore, the modes $\psi_j(\mathbf{x}, f)e^{i2\pi ft}$ are orthogonal in the space–time inner product (2.11), so each mode at each frequency can be viewed as a distinct space–time mode. The space–time correlation tensor can be written as

$$\mathbf{C}(\mathbf{x}, \mathbf{x}', \tau) = \int_{-\infty}^{\infty} \sum_{j=1}^{\infty} \lambda_j(f) \psi_j(\mathbf{x}, f) \psi_j^*(\mathbf{x}', f) e^{i2\pi f\tau} df. \quad (2.22)$$

In summary, for stationary flows, the space–time POD formulation leads to spectral POD modes that each oscillate at a single frequency and optimally represent the second-order space–time flow statistics.

2.3. Spectral versus space-only POD

The essential difference between spectral and space-only POD is that the former yields modes that are coherent in space and time, whereas the later gives modes that are only spatially coherent. First, consider space-only POD. The time dependence of the flow field $\mathbf{q}(\mathbf{x}, t)$ is treated as a stochastic parameter, with different time instances taken to represent different members in an ensemble of spatially dependent fields. Once interpreted in this way, the flow snapshots that make up the ensemble lose any concept of sequential ordering, so the time-dependent evolution of the flow has no impact on the definition of the POD modes. Therefore, the POD modes are impervious to temporal correlation within the data, which is an essential feature of physical coherent structures. Because of this, POD modes do not necessarily represent structures that evolve coherently in space and time. This can be shown explicitly by writing the space–time correlation tensor in terms of space-only POD modes,

$$\mathbf{C}(\mathbf{x}, \mathbf{x}', \tau) = E \left\{ \left(\sum_{j=1}^{\infty} a_j(t) \phi_j(\mathbf{x}) \right) \left(\sum_{k=1}^{\infty} a_k(t + \tau) \phi_k(\mathbf{x}') \right)^* \right\} \quad (2.23a)$$

$$= \sum_{j=1}^{\infty} \sum_{k=1}^{\infty} C_{a_j a_k}^{POD}(\tau) \phi_j(\mathbf{x}) \phi_k^*(\mathbf{x}'), \quad (2.23b)$$

with

$$C_{a_j a_k}^{POD}(\tau) = E\{a_j(t) a_k^*(t + \tau)\}. \quad (2.24)$$

When $\tau = 0$, (2.9) ensures that $C_{a_j a_k}^{POD}(\tau) = \lambda_j \delta_{jk}$ and (2.23) reduces to the spatial correlation tensor. On the other hand, (2.9) is not applicable when $\tau \neq 0$, and, as a result, POD theory does not guarantee any special properties for $C_{a_j a_k}^{POD}(\tau)$; thus, the temporal correlation of two terms in the POD expansion is not known *a priori*. This means that the part of the flow described by a given POD mode is not necessarily correlated with the part of the flow described by the same POD mode at a later time, nor is it necessarily uncorrelated with the part of the flow described by a different mode at a later time. Therefore, contrary to some previous statements, space-only POD modes do not necessarily represent flow structures that evolve coherently. A recent analysis of space-only POD by George (2017) erroneously assumed the expansion coefficients to be uncorrelated at different times.

In contrast, SPOD modes do represent structures that evolve coherently. The space–time correlation tensor is written in terms of SPOD modes in (2.22). This form of the space–time correlation tensor is the result of special properties of the correlations

$$C_{a_j a_k}^{SPOD}(f, f') \triangleq E\{a_j(f) a_k^*(f')\} = \lambda_j \delta_{jk} \delta(f - f'), \quad (2.25)$$

which govern the correlation between the part of the flow described by individual SPOD modes. Specifically, (2.22) can be obtained by inserting (2.25) along with the inverse Fourier transform of (2.19) into the definition of the space–time correlation tensor given by (2.14) and (2.15). The first two terms in the final form of (2.25) follow from (2.20) and ensure that two terms in the SPOD expansion at the same frequency are uncorrelated. The final term is a consequence of the fact that the frequency components from the Fourier transform of a stationary random process are uncorrelated (Lumley 1970; George 1988).

In sum, (2.22) and (2.25) show that the part of the flow described by a particular SPOD mode is perfectly correlated with the part of the flow described by that same mode at all times and entirely uncorrelated with the part of the flow described by all other modes at all times. In other words, each SPOD mode describes a structure that evolves coherently in space and time.

The preceding analysis does not imply that space-only POD modes can never exhibit space–time coherence. For example, Rowley *et al.* (2004) observed that some of the leading space-only POD modes of a compressible cavity flow capture single-frequency Rossiter modes. Rather, our analysis shows that space-only POD modes do not have this property by construction, in contrast to SPOD modes which evolve coherently in space and time by construction.

It is also possible to derive equations relating space-only and spectral POD modes and eigenvalues. Using the fact that the spatial correlation tensor is equivalent to the zero-time-lag space–time correlation tensor, (2.17) implies that

$$\mathbf{C}(\mathbf{x}, \mathbf{x}') = \int_{-\infty}^{\infty} \mathbf{S}(\mathbf{x}, \mathbf{x}', f) df. \quad (2.26)$$

Expanding \mathbf{C} and \mathbf{S} in terms of POD and SPOD modes respectively gives

$$\sum_{j=1}^{\infty} \lambda_j \phi_j(\mathbf{x}) \phi_j(\mathbf{x}')^* = \int_{-\infty}^{\infty} \sum_{k=1}^{\infty} \lambda_k(f) \psi_k(\mathbf{x}, f) \psi_k(\mathbf{x}', f)^* df. \quad (2.27)$$

Applying the operation $\langle \phi_j, \cdot \rangle_x$ to both sides of (2.27) and dividing by λ_j leads to the expression

$$\phi_j(\mathbf{x}) = \int_{-\infty}^{\infty} \sum_{k=1}^{\infty} \frac{\lambda_k(f)}{\lambda_j} c_{jk}(f) \psi_k(\mathbf{x}, f) df, \quad (2.28)$$

where $c_{jk}(f) = \langle \phi_j(\mathbf{x}), \psi_k(\mathbf{x}, f) \rangle_x$. Taking the same inner product again and moving λ_j back to the left-hand side yields

$$\lambda_j = \int_{-\infty}^{\infty} \sum_{k=1}^{\infty} \lambda_k(f) |c_{jk}(f)|^2 df. \quad (2.29)$$

Equations (2.28) and (2.29) show that each space-only POD mode is potentially made up of many SPOD modes. Physically, this means that the spatially coherent structures represented by space-only POD are composed of contributions from spatiotemporal coherent structures at many frequencies. Practically, this is manifested as broadband frequency content within the coefficients $a_j(t)$. This highlights the fact that each space-only POD mode typically represents flow phenomena at many different time scales, which muddies their interpretation. In contrast, SPOD modes decouple flow phenomena at different time scales, which can be helpful for understanding the flow dynamics and deriving non-empirical models.

3. Computing SPOD modes from data

An efficient algorithm for computing space-only POD modes from snapshots of discrete flow data using the method of snapshots (Sirovich 1987) is well known and described in detail in numerous publications (e.g. Rowley & Dawson 2017; Taira *et al.* 2017). Techniques for computing SPOD modes from snapshots of the flow are not as well documented. Here, we outline a procedure similar to the one used by Citriniti & George (2000) and Gordeyev & Thomas (2000) but with an additional simplification that reduces the computational cost in most cases.

Let the vector $\mathbf{q}_k \in \mathbb{R}^N$ represent the instantaneous state of $\mathbf{q}(\mathbf{x}, t)$ at time t_k on a discrete set of points in the spatial domain Ω . The total length N of the vector is equal to the number of grid points times the number of flow variables, since all of these values have been stacked into the vector \mathbf{q}_k , which we call a snapshot of the flow. Now, suppose that these data are available for M equally spaced time instances, $t_{k+1} = t_k + \Delta t$. This data set can be compactly represented by the data matrix

$$\mathbf{Q} = [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_M] \in \mathbb{R}^{N \times M}. \quad (3.1)$$

The cross-spectral density tensor could be naively estimated using the discrete Fourier transform (DFT) of the data matrix \mathbf{Q} . However, it is well known that spectral estimates obtained in this way do not converge as the number of snapshots M is increased. In fact, the uncertainty in the estimate at each frequency is as large as the magnitude of the estimate itself (George, Beuther & Lumley 1978; Bendat & Piersol 2000). To obtain convergent estimates of the spectral densities, it is necessary to appropriately average the spectra over multiple realizations of the flow. This can be accomplished using Welch's (1967) method, which is represented schematically in figure 1. The first step is to partition the data matrix into a set of smaller, possibly overlapping, blocks. Precisely, if we write each block as

$$\mathbf{Q}^{(n)} = [\mathbf{q}_1^{(n)}, \mathbf{q}_2^{(n)}, \dots, \mathbf{q}_{N_f}^{(n)}] \in \mathbb{R}^{N \times N_f}, \quad (3.2)$$

then the k th entry in the n th block is $\mathbf{q}_k^{(n)} = \mathbf{q}_{k+(n-1)(N_f-N_o)}$, where N_f is the number of snapshots in each block, N_o is the number of snapshots by which the blocks overlap and N_b is the total number of blocks. By the ergodicity hypothesis, each of these blocks can be regarded as a member of an ensemble of realizations of the flow.

Next, the DFT is computed for each block,

$$\hat{\mathbf{Q}}^{(n)} = [\hat{\mathbf{q}}_1^{(n)}, \hat{\mathbf{q}}_2^{(n)}, \dots, \hat{\mathbf{q}}_{N_f}^{(n)}], \quad (3.3)$$

with

$$\hat{\mathbf{q}}_k^{(n)} = \frac{1}{\sqrt{N_f}} \sum_{j=1}^{N_f} w_j \mathbf{q}_j^{(n)} e^{-i2\pi(k-1)[(j-1)/N_f]} \quad (3.4)$$

for $k = 1, \dots, N_f$ and $n = 1, \dots, N_b$. The scalar weights w_j are nodal values of a window function that can be used to reduce spectral leakage due to non-periodicity of the data in each block (e.g. Heinzel, Rüdiger & Schilling 2002). We have included the $1/\sqrt{N_f}$ factor to make the discrete transform unitary for a rectangular window ($w_j = 1$ for all j), which will be convenient later in § 4. Here, $\hat{\mathbf{q}}_k^{(n)}$ is the Fourier component




FIGURE 1. (Colour online) Schematic depiction of Welch's method for estimating SPOD modes. A detailed description of each step is given in the text.

at frequency f_k in the n th block and the resolved frequencies are

$$f_k = \begin{cases} \frac{k-1}{N_f \Delta t} & \text{for } k \leq N_f/2, \\ \frac{k-1-N_f}{N_f \Delta t} & \text{for } k > N_f/2. \end{cases} \quad (3.5)$$

The cross-spectral density tensor $\mathbf{S}(\mathbf{x}, \mathbf{x}', f)$ can be estimated at frequency f_k by the average

$$\mathbf{S}_{f_k} = \frac{\Delta t}{s N_b} \sum_{n=1}^{N_b} \hat{\mathbf{q}}_k^{(n)} (\hat{\mathbf{q}}_k^{(n)})^*, \quad (3.6)$$

where $s = \sum_{j=1}^{N_f} w_j^2$. This can be written compactly by arranging the Fourier coefficients at frequency f_k from each block into the new data matrix

$$\hat{\mathbf{Q}}_{f_k} = \sqrt{\kappa} [\hat{\mathbf{q}}_k^{(1)}, \hat{\mathbf{q}}_k^{(2)}, \dots, \hat{\mathbf{q}}_k^{(N_b)}] \in \mathbb{R}^{N \times N_b}, \quad (3.7)$$

where $\kappa = \Delta t / (s N_b)$. Then, the estimated cross-spectral density tensor at frequency f_k can be written as

$$\mathbf{S}_{f_k} = \hat{\mathbf{Q}}_{f_k} \hat{\mathbf{Q}}_{f_k}^*. \quad (3.8)$$

This estimate converges as the number of blocks N_b and the number of snapshots in each block N_f are increased together (Welch 1967; Bendat & Piersol 2000).

Using this estimate, the infinite-dimensional SPOD eigenvalue problem (2.18) reduces to an $N \times N$ matrix eigenvalue problem,

$$\mathbf{S}_{f_k} \mathbf{W} \Psi_{f_k} = \Psi_{f_k} \Lambda_{f_k} \quad (3.9)$$

at each frequency. Here, the spatial inner product (2.4) is approximated as $\langle \hat{\mathbf{q}}_1, \hat{\mathbf{q}}_2 \rangle_x = \hat{\mathbf{q}}_2^* \mathbf{W} \hat{\mathbf{q}}_1$; the positive-definite Hermitian matrix $\mathbf{W} \in \mathbb{C}^{N \times N}$ accounts for both the

weight $\mathbf{W}(\mathbf{x})$ and the numerical quadrature of the integral on the discrete grid. The approximate SPOD modes are given by the columns of Ψ_{f_k} and are ranked according to their corresponding eigenvalues given by the diagonal matrix Λ_{f_k} . Note that at most N_b non-zero eigenvalues can be obtained. The approximate modes mimic the properties of the continuous modes. For example, they are discretely orthogonal, $\Psi_{f_k}^* \mathbf{W} \Psi_{f_k} = \mathbf{I}$, and the estimated cross-spectral density tensor can be expanded as $\mathbf{S}_{f_k} = \Psi_{f_k} \Lambda_{f_k} \Psi_{f_k}^*$.

In practice, the number of blocks N_b is typically much smaller than the discretized problem size N . Using the definition of \mathbf{S}_{f_k} from (3.8), it is possible to show that the $N_b \times N_b$ eigenvalue problem

$$\hat{\mathbf{Q}}_{f_k}^* \mathbf{W} \hat{\mathbf{Q}}_{f_k} \Theta_{f_k} = \Theta_{f_k} \tilde{\Lambda}_{f_k} \quad (3.10)$$

supports the same non-zero eigenvalues as (3.9). The eigenvectors corresponding to these non-zero eigenvalues can be exactly recovered as

$$\tilde{\Psi}_{f_k} = \hat{\mathbf{Q}}_{f_k} \Theta_{f_k} \tilde{\Lambda}_{f_k}^{-1/2}. \quad (3.11)$$

The complete procedure for computing SPOD modes from data snapshots is outlined in the following algorithm. Variables that are assigned using the ‘ \leftarrow ’ operator can be deleted or overwritten after each iteration in their respective loop to reduce memory usage. A MATLAB implementation is available at https://github.com/SpectralPOD/spod_matlab. We also note that Schmidt (2017) recently formulated a streaming version of the algorithm that can reduce computational cost for large data sets.

Algorithm (SPOD)

1. For each data block $n = 1, 2, \dots, N_b$:

- (a) assemble the data matrix

$$\mathbf{Q}^{(n)} \leftarrow [\mathbf{q}_{1+(n-1)(N_f-N_o)}, \mathbf{q}_{2+(n-1)(N_f-N_o)}, \dots, \mathbf{q}_{N_f+(n-1)(N_o-N_o)}];$$

- (b) using a (windowed) fast Fourier transform, calculate and store the row-wise DFT

$$\hat{\mathbf{Q}}^{(n)} = \text{FFT}(\mathbf{Q}^{(n)}) = [\hat{\mathbf{q}}_1^{(n)}, \hat{\mathbf{q}}_2^{(n)}, \dots, \hat{\mathbf{q}}_{N_o}^{(n)}].$$

The column $\hat{\mathbf{q}}_k^{(n)}$ contains the n th realization of the Fourier mode at the k th discrete frequency f_k .

2. For each frequency $k = 1, 2, \dots, N_f$ (or some subset of interest):

- (a) assemble the matrix $\hat{\mathbf{Q}}_{f_k} \leftarrow \sqrt{\kappa} [\hat{\mathbf{q}}_k^{(1)}, \hat{\mathbf{q}}_k^{(2)}, \dots, \hat{\mathbf{q}}_k^{(N_b)}]$ of Fourier realizations from the k th column of each $\hat{\mathbf{Q}}^{(n)}$;

- (b) calculate the matrix $\mathbf{M}_{f_k} \leftarrow \hat{\mathbf{Q}}_{f_k}^* \mathbf{W} \hat{\mathbf{Q}}_{f_k}$;

- (c) compute the eigenvalue decomposition $\mathbf{M}_{f_k} = \Theta_{f_k} \tilde{\Lambda}_{f_k} \Theta_{f_k}^*$;

- (d) obtain and store the SPOD modes $\tilde{\Psi}_{f_k} = \hat{\mathbf{Q}}_{f_k} \Theta_{f_k} \tilde{\Lambda}_{f_k}^{-1/2}$ and modal energies $\tilde{\Lambda}_{f_k}$ for the k th discrete frequency.

4. Spectral POD and DMD

In this section, we investigate the relationship between SPOD and DMD, and show that SPOD can be understood as an optimal form of DMD for statistically stationary turbulent flows. Dynamic mode decomposition was developed as an alternative to POD for identifying coherent structures from flow data (Schmid & Sesterhenn 2008; Rowley *et al.* 2009; Schmid 2010). The method approximates the eigenmodes of a linear operator that maps the state of the flow from one time instant to the next. Since the operator is linear, the temporal evolution of each mode is described by a single frequency and growth/decay rate, and the modes are in general spatially non-orthogonal. Just as space-only POD can be described as a spatial orthogonalization of the flow data, DMD can be understood as a temporal orthogonalization of the data (Schmid 2010).

4.1. DMD and the DFT

For zero-mean data that are uniformly sampled in time, Chen, Tu & Rowley (2012) showed that DMD is formally equivalent to the DFT (the flow snapshots must also be linearly independent for this to hold, which is true in most applications). As a result, each DMD mode has zero growth/decay rate. This formal DMD–DFT equivalence does not hold for data with a non-zero mean, but in practice the physically relevant DMD modes tend to have nearly zero growth/decay rates for stationary flows (e.g. Rowley *et al.* 2009; Chen *et al.* 2012; Schmid, Violato & Scarano 2012; Semeraro, Bellani & Lundell 2012), which is logical since stationary flows are persistent by definition.

This tendency is explained more rigorously by the connection between DMD and Koopman operator theory (Mezić 2005; Rowley *et al.* 2009). The Koopman operator is an infinite-dimensional linear operator that describes the evolution of scalar observables of a nonlinear dynamical system on a finite manifold. The eigenvectors of the Koopman operator can be used to decompose vector-valued observables (equivalent to our \mathbf{q}) into modes that evolve with a single frequency and growth/decay rate. Mezić (2005) showed that for any dynamical system with a Borel probability measure, the growth/decay rate is zero and Koopman modes are equivalent to Fourier modes. Stationary flows possess an ergodic measure by definition, so their Koopman modes are simply Fourier modes.

Rowley *et al.* (2009) showed that DMD modes approximate Koopman modes when the flow snapshots used to compute the modes are linearly independent, and Tu *et al.* (2014) showed that this holds under a slightly weaker condition on the data termed linear consistency. In light of this, it is not surprising that DMD modes tend to be similar to DFT modes for stationary flows. Moreover, deviations of the DMD modes from Fourier modes can be viewed as artefacts of the DMD approximation of the underlying Koopman operator. This suggests that, contrary to prevailing wisdom, it is advantageous to subtract the mean when applying DMD to stationary flows to ensure that the DMD modes mimic the zero-growth-rate property of the underlying Koopman modes.

4.2. An ensemble DMD problem for stationary flows

To conceptually relate SPOD and DMD, imagine that we have an ensemble of N_e data sets, each representing a stochastic realization of the same stationary flow. There are at least two approaches one could adopt for applying DMD to this problem.

The first would be to simply perform separate DMD computations for each realization of the flow. If we agree to subtract the mean to ensure zero growth rates as described previously, each DMD mode will exactly correspond to a DFT mode, but in general the mode at a given frequency will be different for each realization.

A more sophisticated approach for applying DMD to this problem would be to use the approach of Tu *et al.* (2014) to combine multiple flow realization into a single DMD calculation. Their variant of DMD, which they call ‘exact DMD’, is defined in terms of the operator

$$\mathbf{A} \triangleq \mathbf{Y}\mathbf{X}^+, \quad (4.1)$$

where \mathbf{X}^+ is the pseudo-inverse of \mathbf{X} and the columns of the matrices \mathbf{X} and \mathbf{Y} are input–output data pairs that are related by a linear operator that is to be approximated by the matrix \mathbf{A} . The DMD modes and eigenvalues are then given by the eigenvectors and eigenvalues of \mathbf{A} respectively.

For a standard application in which the flow data consist of sequential snapshots of the flow, the input and output matrices are

$$\mathbf{X} \triangleq [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_{M-1}] = \mathbf{Q}\mathbf{T}_X, \quad (4.2a)$$

$$\mathbf{Y} \triangleq [\mathbf{q}_2, \mathbf{q}_3, \dots, \mathbf{q}_M] = \mathbf{Q}\mathbf{T}_Y, \quad (4.2b)$$

where each $\mathbf{q}_j \in \mathbb{R}^N$ is a snapshot of the flow as defined in § 3 and $\mathbf{Q} \in \mathbb{R}^{N \times M}$ is the data matrix given by (3.1). The matrices $\mathbf{T}_X, \mathbf{T}_Y \in \mathbb{R}^{M \times M-1}$ select the appropriate columns of \mathbf{Q} and are given by

$$\mathbf{T}_X \triangleq \begin{bmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & 1 \\ 0 & 0 & \cdots & 0 \end{bmatrix}, \quad \mathbf{T}_Y \triangleq \begin{bmatrix} 0 & 0 & \cdots & 0 \\ 1 & 0 & \cdots & 0 \\ 0 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & 1 \end{bmatrix}. \quad (4.3a,b)$$

As described by Tu *et al.* (2014), multiple realizations of a flow can be accommodated within the exact DMD framework by arranging the realizations together into ensemble input and output matrices,

$$\mathbf{X} \triangleq [\mathbf{Q}^{(1)}\mathbf{T}_X, \dots, \mathbf{Q}^{(N_e)}\mathbf{T}_X], \quad (4.4a)$$

$$\mathbf{Y} \triangleq [\mathbf{Q}^{(1)}\mathbf{T}_Y, \dots, \mathbf{Q}^{(N_e)}\mathbf{T}_Y], \quad (4.4b)$$

where each $\mathbf{Q}^{(n)} \in \mathbb{R}^{N \times M}$ contains snapshots from the n th realization of the flow, as in (3.2).

The DMD/DFT equivalence for zero-mean data proven by Chen *et al.* (2012) was derived within the context of the original Arnoldi-based formulation of DMD given by Schmid (2010) and was restricted to the case of sequential linearly independent snapshots. Accordingly, it does not apply to the ensemble DMD problem defined by (4.1) and (4.4).

In appendix B, we prove that the DMD modes obtained from this ensemble formulation are precisely the DFT modes of each realization of the flow if each realization has zero mean (i.e. we have subtracted the mean in line with our earlier

arguments for stationary flows) and the ensemble input/output data matrices are linearly consistent. Under these conditions, (4.1) and (4.4) reduce to

$$\mathbf{A} \begin{bmatrix} \tilde{\mathbf{Q}}^{(1)}, \dots, \tilde{\mathbf{Q}}^{(N_e)} \end{bmatrix} = \begin{bmatrix} \tilde{\mathbf{Q}}^{(1)}, \dots, \tilde{\mathbf{Q}}^{(N_e)} \end{bmatrix} \begin{bmatrix} \mathbf{\Lambda}_{\text{DMD}} & & \\ & \ddots & \\ & & \mathbf{\Lambda}_{\text{DMD}} \end{bmatrix}, \quad (4.5)$$

where the columns of $\tilde{\mathbf{Q}}^{(n)}$ contain the DFT modes of $\mathbf{Q}^{(n)}$ excluding the mean component,

$$\mathbf{\Lambda}_{\text{DMD}} = \begin{bmatrix} z & & & \\ & z^2 & & \\ & & \ddots & \\ & & & z^{M-1} \end{bmatrix}, \quad (4.6)$$

and $z = e^{-i2\pi/M}$ is the primitive M th root of unity. Therefore, each DFT mode $\hat{\mathbf{q}}_k^{(n)}$, $k = 2, \dots, M$, is an eigenvector of \mathbf{A} with eigenvalue z^{k-1} . These eigenvalues have unit magnitude, so the growth/decay rate is zero, as expected. The frequency corresponding to each DFT mode is recovered from the eigenvalue as $f_k = \text{Re}[\ln(z^{k-1})/(-i2\pi\Delta t)]$, and it is easy to verify that these frequencies match those given by (3.5). Therefore, each DFT mode from each realization of the flow is a DMD mode that oscillates at the corresponding DFT frequency. Since each eigenvalue is repeated N_e times, any linear combination of the DFT modes at a given frequency is also a DMD mode.

4.3. DMD and SPOD

Overall, for either the simple approach in which a separate \mathbf{A} operator is defined for each flow realization or the more sophisticated approach in which multiple realizations are incorporated into a single ensemble DMD computation, the end result is a set of DMD modes at each frequency consisting of the DFT mode of each realization (or a linear combination thereof). Because of the random nature of turbulent flows and the uncertainty inherent to the DFT, the mode for a given frequency obtained from each flow realization will in general be different. This statistical variability may seem undesirable, but in fact it exposes an important characteristic of turbulent flows – the behaviour at a given frequency cannot be described by a single deterministic mode, i.e. the cross-spectral density tensor is not rank one. In light of this, how can useful information be extracted from this set of differing DMD modes? A sensible approach is to search for functions that best represent the ensemble of DMD modes at each frequency; these are given precisely by the SPOD modes at that frequency. In other words, SPOD modes provide the optimal basis, as defined by (2.12), for describing the variability within the ensemble of DMD modes. Said the other way around, each DMD mode in the ensemble at a particular frequency represents one of the possible ways in which the coherent structures represented by the SPOD modes can coexist in one particular realization of the turbulent flow.

The preceding analysis of the ensemble DMD problem enables us to make a still stronger connection between DMD and SPOD for stationary flows. As already mentioned, any linear combination of the set of DFT modes at a given frequency is also a DMD mode. Recall from (3.11) that each SPOD mode can be written as a linear combination of the DFT modes used to estimate the cross-spectral density tensor. Therefore, each SPOD mode is itself a DMD mode. Specifically, SPOD modes are given by the particular linear combination, or weighted average, of DFT/DMD

modes that yields uncorrelated modes that optimally capture the flow energy, as defined by (2.12). Therefore, SPOD modes are ensemble DMD modes that have been optimally averaged to provide the best possible representation of the second-order space–time flow statistics. This means that SPOD modes are dynamically significant in the same sense as standard DMD modes, but at the same time optimally describe the random nature of turbulent flows. This connection between SPOD and DMD suggests that SPOD could be directly related to the concept of the stochastic Koopman operator (Mezić 2005), but this remains to be explored in detail.

5. Spectral POD and resolvent analysis

In this section, a connection is made between SPOD and resolvent analysis for stationary flows. Resolvent analysis (Jovanović & Bamieh 2005; Bagheri *et al.* 2009; McKeon & Sharma 2010; Sipp *et al.* 2010) identifies modes that optimally describe the linear growth/amplification mechanisms within the Navier–Stokes equations, and is based on analysis of the linearized Navier–Stokes equations rather than time-resolved data. This approach has been previously related to other modal decomposition techniques. Dergham, Sipp & Robinet (2013) showed that resolvent modes can be used to approximate the controllability Gramian of a linear time-invariant system, the eigenvectors of which are equivalent to space-only POD modes if the system is forced by white noise (Rowley 2005; Bagheri *et al.* 2009). Sharma *et al.* (2016) recently suggested a connection between DMD/Koopman operator theory and resolvent analysis. Specifically, they showed that the resolvent operator relates Koopman modes of the input and output for a flow on an attractor.

The connection that we make between SPOD and resolvent modes is based on a statistical perspective on the resolvent-mode expansion of stationary turbulent flows. In the following sections, we provide a brief development of the basic resolvent theory to introduce our notation and terminology, make a connection between SPOD and resolvent modes, and discuss the implications of this connection.

5.1. Resolvent analysis

We begin with nonlinear flow equations of the form

$$\frac{\partial \chi}{\partial t} = \mathcal{F}(\chi), \quad (5.1)$$

where χ is a state vector of flow variables. The compressible Navier–Stokes equations are naturally written in the form of (5.1), whereas the incompressible Navier–Stokes equations can be written in this form by eliminating the pressure by projecting the velocity field onto a divergence-free basis. Substituting the standard Reynolds decomposition

$$\chi(x, t) = \bar{\chi}(x) + \chi'(x, t) \quad (5.2)$$

into (5.1) and isolating the terms that are linear in χ' yields an equation of the form

$$\frac{\partial \chi'}{\partial t} - \mathcal{A}(\bar{\chi})\chi' = \mathbf{h}(\bar{\chi}, \chi'), \quad (5.3)$$

where

$$\mathcal{A}(\bar{\chi}) = \frac{\partial \mathcal{F}}{\partial \chi}(\bar{\chi}) \quad (5.4)$$

is the linearized flow operator and \mathbf{h} contains the remaining nonlinear terms as well as any exogenous inputs such as incoming fluctuations at boundaries or environmental disturbances.

To make the ensuing analysis as flexible as possible, it is useful to introduce two new variables: an input variable $\boldsymbol{\eta}$ and an output variable \mathbf{y} , defined by the relations

$$\mathbf{h}(\mathbf{x}, t) = \mathcal{B}\boldsymbol{\eta}(\mathbf{x}, t) \quad (5.5)$$

and

$$\mathbf{y}(\mathbf{x}, t) = \mathcal{C}\boldsymbol{\chi}'(\mathbf{x}, t) \quad (5.6)$$

respectively. The linear operator $\mathcal{C} = \mathcal{C}(\mathbf{x})$ can be used to select certain flow variables and/or spatial regions that are of particular interest, or in general any linear function of the state perturbation. Similarly, the linear operator $\mathcal{B} = \mathcal{B}(\mathbf{x})$ limits the allowable form of \mathbf{h} , which can be useful for enforcing known properties of the nonlinearity in certain flows, e.g. that the incompressible continuity equation is linear so that the corresponding component of \mathbf{h} must be zero. Since the flow is stationary, all quantities can be Fourier decomposed, and (5.3), (5.5) and (5.6) can be written in the frequency domain as

$$(i2\pi f\mathcal{I} - \mathcal{A})\hat{\boldsymbol{\chi}} = \hat{\mathbf{h}} = \mathcal{B}\hat{\boldsymbol{\eta}} \quad (5.7)$$

and

$$\hat{\mathbf{y}} = \mathcal{C}\hat{\boldsymbol{\chi}}. \quad (5.8)$$

Following McKeon & Sharma (2010), the basic objective of resolvent analysis is to find pairs of inputs and outputs at each frequency that are optimal in terms of their linear gain

$$\sigma^2 = \frac{\langle \hat{\mathbf{y}}, \hat{\mathbf{y}} \rangle_{\mathbf{y}}}{\langle \hat{\boldsymbol{\eta}}, \hat{\boldsymbol{\eta}} \rangle_{\boldsymbol{\eta}}}, \quad (5.9)$$

where

$$\langle \hat{\mathbf{y}}_1, \hat{\mathbf{y}}_2 \rangle_{\mathbf{y}} = \int_{\Omega} \hat{\mathbf{y}}_2^*(\mathbf{x}, f) \mathbf{W}_{\mathbf{y}}(\mathbf{x}) \hat{\mathbf{y}}_1(\mathbf{x}, f) \, d\mathbf{x} \quad (5.10)$$

and

$$\langle \hat{\boldsymbol{\eta}}_1, \hat{\boldsymbol{\eta}}_2 \rangle_{\boldsymbol{\eta}} = \int_{\Omega} \hat{\boldsymbol{\eta}}_2^*(\mathbf{x}, f) \mathbf{W}_{\boldsymbol{\eta}}(\mathbf{x}) \hat{\boldsymbol{\eta}}_1(\mathbf{x}, f) \, d\mathbf{x} \quad (5.11)$$

are inner products on the output and input spaces respectively.

Using (5.7) and (5.8), the input and output are related as

$$\hat{\mathbf{y}}(\mathbf{x}, f) = \mathcal{R}(\mathbf{x}, f) \hat{\boldsymbol{\eta}}(\mathbf{x}, f), \quad (5.12)$$

where

$$\mathcal{R}(\mathbf{x}, f) = \mathcal{C}(\mathbf{x})(i2\pi f\mathcal{I} - \mathcal{A}(\mathbf{x}))^{-1}\mathcal{B}(\mathbf{x}) \quad (5.13)$$

is termed the resolvent operator. Note that $(i2\pi f\mathcal{I} - \mathcal{A})$ is invertible for all f if \mathcal{A} is stable; i.e. the real part of all of its eigenvalues is negative.

The optimal inputs and outputs are defined by the Schmidt decomposition of the resolvent operator

$$\mathcal{R}(\mathbf{x}, f) = \sum_{j=1}^{\infty} \sigma_j(f) \hat{\mathbf{u}}_j(\mathbf{x}, f) \otimes \hat{\mathbf{v}}_j(\mathbf{x}, f), \quad (5.14)$$

defined in terms of the inner products (5.10) and (5.11). The input modes $\hat{\mathbf{v}}_j$ and output modes $\hat{\mathbf{u}}_j$ provide complete bases for their respective spaces and are orthogonal in their respective inner products (i.e. $\langle \hat{\mathbf{u}}_j, \hat{\mathbf{u}}_k \rangle_{\mathbf{y}} = \langle \hat{\mathbf{v}}_j, \hat{\mathbf{v}}_k \rangle_{\boldsymbol{\eta}} = \delta_{jk}$). The modes are ordered

by their singular value σ_j , and the gain between $\hat{\mathbf{v}}_j$ and $\hat{\mathbf{u}}_j$ is σ_j^2 . The resolvent operator is said to be low rank if the magnitude of the singular values falls off rapidly with increasing j .

Since the bases are complete, the output can be expanded as

$$\hat{\mathbf{y}}(\mathbf{x}, f) = \sum_{j=1}^{\infty} \hat{\mathbf{u}}_j(\mathbf{x}, f) \sigma_j(f) \beta_j(f), \quad (5.15)$$

where

$$\beta_j(f) = \langle \hat{\boldsymbol{\eta}}(\mathbf{x}, f), \hat{\mathbf{v}}_j(\mathbf{x}, f) \rangle_{\eta} \quad (5.16)$$

is the projection of the forcing onto the j th input mode (McKeon & Sharma 2010). In some flows, reduced-order models have been obtained by retaining a limited number of terms in the expansion when the resolvent operator is low rank (e.g. Beneddine *et al.* 2016; Gómez *et al.* 2016b).

5.2. Spectral densities in terms of resolvent modes

A connection between resolvent analysis and SPOD can be obtained by writing the cross-spectral density tensor in terms of resolvent modes. To do so, it is convenient to write the cross-spectral density tensor as

$$\mathbf{S}_{yy}(\mathbf{x}, \mathbf{x}', f) = E\{\hat{\mathbf{y}}(\mathbf{x}, f) \hat{\mathbf{y}}^*(\mathbf{x}', f)\}, \quad (5.17)$$

which can be proven to be equivalent to the previous definition given by (2.16) (Bendat & Piersol 2000). Inserting the resolvent-mode expansion (5.15) leads to the expression

$$\mathbf{S}_{yy}(\mathbf{x}, \mathbf{x}', f) = \sum_{j=1}^{\infty} \sum_{k=1}^{\infty} \hat{\mathbf{u}}_j(\mathbf{x}, f) \hat{\mathbf{u}}_k^*(\mathbf{x}', f) \sigma_j(f) \sigma_k(f) S_{\beta_j \beta_k}(f), \quad (5.18)$$

where

$$S_{\beta_j \beta_k}(f) = E\{\beta_j(f) \beta_k^*(f)\} \quad (5.19)$$

is the scalar-valued cross-spectral density between the j th and k th expansion coefficients. In obtaining (5.19), the output resolvent modes and singular values were moved outside of the expectation operator, which is permitted because they are deterministic quantities, depending only on the resolvent operator and the inner products. On the other hand, the expansion coefficients depend on the forcing term $\hat{\boldsymbol{\eta}}$, which is stochastic due to the random nature of turbulent flows, so the expansion coefficients must be described by their cross-spectral density. This fact is central to the remainder of this section.

A connection between SPOD and resolvent analysis is obtained by choosing the SPOD quantity of interest \mathbf{q} and the resolvent output \mathbf{y} to be the same, with the same inner products, $\langle \cdot, \cdot \rangle_x = \langle \cdot, \cdot \rangle_y$. Then, the cross-spectral density tensors \mathbf{S} and \mathbf{S}_{yy} are identical and the SPOD expansion (2.21) and resolvent-mode expansion (5.18) can be equated,

$$\mathbf{S}_{yy}(\mathbf{x}, \mathbf{x}', f) = \sum_{j=1}^{\infty} \lambda_j(f) \boldsymbol{\psi}_j(\mathbf{x}, f) \boldsymbol{\psi}_j^*(\mathbf{x}', f) \quad (5.20a)$$

$$= \sum_{j=1}^{\infty} \sum_{k=1}^{\infty} \hat{\mathbf{u}}_j(\mathbf{x}, f) \hat{\mathbf{u}}_k^*(\mathbf{x}', f) \sigma_j(f) \sigma_k(f) S_{\beta_j \beta_k}(f). \quad (5.20b)$$

The implications of (5.20) will be explored in the following two subsections.

5.3. Special case: uncorrelated expansion coefficients

The resolvent-mode expansion of the cross-spectral density tensor can be simplified in the special case in which the β_j expansion coefficients are uncorrelated from one another, which is expressed as $S_{\beta_j\beta_k}(f) = \mu_j(f)\delta_{jk}$. Substituting this into (5.20b) leads to the simplified expression

$$\mathbf{S}_{yy}(\mathbf{x}, \mathbf{x}', f) = \sum_{j=1}^{\infty} \boldsymbol{\psi}_j(\mathbf{x}, f) \boldsymbol{\psi}_j^*(\mathbf{x}', f) \lambda_j(f) \quad (5.21a)$$

$$= \sum_{j=1}^{\infty} \hat{\mathbf{u}}_j(\mathbf{x}, f) \hat{\mathbf{u}}_j^*(\mathbf{x}', f) \sigma_j^2(f) \mu_j(f). \quad (5.21b)$$

Since the diagonalization of a normal operator via a basis that is orthogonal in a given inner product is unique, (5.21) shows that the sets of SPOD and resolvent modes are identical. Specifically, the resolvent mode with maximum $\sigma_j^2 \mu_j$ corresponds to the first SPOD mode, and so on. If $\mu_j = 1$ for every j , then the ordering of the two sets of modes is the same and $\sigma_j^2(f) = \lambda_j(f)$, $\hat{\mathbf{u}}_j(\mathbf{x}, f) = \boldsymbol{\psi}_j(\mathbf{x}, f)$.

The conditions for which the expansion coefficients are uncorrelated can be elucidated by manipulating the definition of $S_{\beta_j\beta_k}$ using properties of inner products,

$$S_{\beta_j\beta_k} = E\{\langle \hat{\boldsymbol{\eta}}, \hat{\mathbf{v}}_j \rangle_{\eta} \langle \hat{\boldsymbol{\eta}}, \hat{\mathbf{v}}_k \rangle_{\eta}^* \} = E\{\langle \langle \hat{\boldsymbol{\eta}}, \hat{\mathbf{v}}_j \rangle_{\eta} \hat{\boldsymbol{\eta}}^*, \hat{\mathbf{v}}_k \rangle_{\eta}^* \} = \langle \langle \mathbf{S}_{\eta\eta}, \hat{\mathbf{v}}_j \rangle_{\eta}^*, \hat{\mathbf{v}}_k \rangle_{\eta}^*, \quad (5.22)$$

where $\mathbf{S}_{\eta\eta}(\mathbf{x}, \mathbf{x}', f) = E\{\hat{\boldsymbol{\eta}}(\mathbf{x}, f) \hat{\boldsymbol{\eta}}^*(\mathbf{x}', f)\}$ is the cross-spectral density tensor of the input η . Recalling the orthogonality of the resolvent output modes, the last form of (5.22) shows that $S_{\beta_j\beta_k} = \mu_j \delta_{jk}$ only if $\langle \mathbf{S}_{\eta\eta}(\mathbf{x}, \mathbf{x}', f), \hat{\mathbf{v}}_j(\mathbf{x}', f) \rangle_{\eta}^* = \mu_j \hat{\mathbf{v}}_j(\mathbf{x}, f)$. Using the definition of the inner product, this can be written as

$$\int_{\Omega} \mathbf{S}_{\eta\eta}(\mathbf{x}, \mathbf{x}', f') \mathbf{W}_{\eta}(\mathbf{x}') \mathbf{v}_j(\mathbf{x}', f') d\mathbf{x}' = \mu(f') \mathbf{v}_j(\mathbf{x}, f'). \quad (5.23)$$

This is precisely the eigenvalue problem that defines the SPOD modes of the input. Thus, uncorrelated expansion coefficients imply that the resolvent input modes are the same as the SPOD modes of the input (if there are repeated eigenvalues/singular values, then the corresponding modes must span the same subspace). Conversely, if the SPOD modes of the input are the same as the input resolvent modes, inserting the SPOD expansion of $\mathbf{S}_{\eta\eta}$ into (5.22) shows that the expansion coefficients are uncorrelated. Therefore, the expansion coefficients are uncorrelated if and only if the resolvent input modes correspond exactly with the SPOD modes of the input.

On inspection of (5.23), the stronger condition that $\mu_j = 1$ for every j requires that

$$\mathbf{S}_{\eta\eta}(\mathbf{x}, \mathbf{x}', f) \mathbf{W}_{\eta}(\mathbf{x}') = \mathbf{I} \delta(\mathbf{x} - \mathbf{x}'). \quad (5.24)$$

When $\mathbf{W}_{\eta} = \mathbf{I}$, this condition takes on physical significance: $\mathbf{S}_{\eta\eta}(\mathbf{x}, \mathbf{x}', f) = \mathbf{I} \delta(\mathbf{x} - \mathbf{x}')$ corresponds to an input that is completely uncorrelated in space and time and has unit amplitude everywhere, i.e. unit-amplitude white noise. While the nonlinear forcing terms in any real flow are unlikely to be white, this approximation has been shown to be reasonable in some flows and is frequently used for the construction of low-order models (Farrell & Ioannou 1993, 1996, 2001), including resolvent-based models (Jovanović & Bamieh 2001; Bagheri *et al.* 2009; Sipp *et al.* 2010; Moarref & Jovanović 2012; Dergham *et al.* 2013). In these models, then, resolvent modes can be understood as approximations of SPOD modes. The equivalence of SPOD and resolvent modes in the case of white-noise forcing has been recently pointed out by multiple authors working in the area of jet-noise modelling (Towne *et al.* 2015; Semeraro *et al.* 2016a; Towne, Brès & Lele 2016).

5.4. General case: flow reconstruction with correlated expansion coefficients

The nonlinear forcing terms in real flows are not white (Towne, Brès & Lele 2017; Zare, Jovanović & Georgiou 2017), which leads to correlated resolvent-mode expansion coefficients. In this case, a more general relationship between SPOD and resolvent modes can be derived that provides insight into how the resolvent-mode expansion coefficients can be chosen to optimally reconstruct the flow.

A number of strategies for choosing the resolvent-mode expansion coefficients have been proposed and employed in the last few years. Perhaps the most straightforward approach, used for example by Jeun, Nichols & Jovanović (2016), is to use simulation data to compute η , estimate $\hat{\eta}$ using a single DFT of the time series, and calculate the expansion coefficients using their definition in (5.16). However, this method suffers from the uncertainty inherent in the DFT, as discussed earlier. Methods using converged statistics of the flow must be used to eliminate this uncertainty.

Moarref *et al.* (2014) formulated a convex optimization problem to find a vector of expansion coefficients that optimally reproduce experimentally measured streamwise energy spectra, averaged in the other directions, for a turbulent channel flow. The coefficients that were optimal in the sense they defined led to qualitative agreement with the statistics, but notably the error did not tend towards zero as the number of resolvent modes was increased.

Gómez *et al.* (2016a) derived an equation for the power spectral densities of the expansion coefficients, which define their magnitude. These are obtained by projecting the time-varying forcing term η onto the resolvent input basis for a range of frequencies, and computing the power spectral densities of these terms. In a separate paper (Gómez *et al.* 2016b), the same authors suggested a somewhat different method using a time series of data at one location in the flow. In this approach, the resolvent expansion is inverse Fourier transformed back into the time domain and the coefficients for all frequencies are simultaneously determined by matching the time series in a least-squares sense.

Beneddine *et al.* (2016) used two different approaches for finding the coefficient for (only) the leading resolvent mode at each frequency for the flow over a backward-facing step. In the first approach, the amplitudes of the expansion coefficients were chosen so that the rank-one resolvent-mode reconstruction matched the power spectral density of the spanwise-averaged streamwise velocity at a single location in the flow. In the second approach, the expansion coefficient for the leading mode at each frequency was chosen so that the rank-one resolvent expansion matched the leading SPOD mode, again at a single location. They found this second method to provide improved estimates of the power spectral density at other spatial locations in the flow. Their results were further improved by attempting to match the SPOD mode at several locations using a least-squares approach.

All of these methods assume that each expansion coefficient can be described by a single deterministic amplitude and phase. In contrast, (5.18) shows that the cross-spectral densities of the expansion coefficients are required to properly reconstruct the second-order statistics of the flow. Notably, this includes the power spectral density, which is obtained by setting $\mathbf{x}' = \mathbf{x}$. Treating the expansion coefficients as deterministic quantities imposes a fundamental limitation on the quality of the flow approximation that can be achieved by the resolvent model. For clarity, we will use the symbol b_j in place of β_j to denote these deterministic expansion coefficients. Regardless of how the b_j values are chosen, the approximation of $S_{\beta_j \beta_k}$ implied by these coefficients is $S_{b_j b_k} = b_j b_k^*$. Making this substitution in (5.18), the resulting cross-spectral density tensor can

be factored into the form

$$\mathbf{s}_{yy}(\mathbf{x}, \mathbf{x}', f) = \left(\sum_{j=1}^{\infty} \hat{\mathbf{u}}_j(\mathbf{x}, f) \sigma_j(f) b_j(f) \right) \left(\sum_{k=1}^{\infty} \hat{\mathbf{u}}_k(\mathbf{x}', f) \sigma_k(f) b_k(f) \right)^*. \quad (5.25)$$

Since the right-hand side is the product of two vectors, this approximation of the cross-spectral density tensor has rank equal to one, whereas the correct cross-spectral density tensor has rank equal to the number of non-zero SPOD eigenvalues at each frequency. This limits the quality of the flow approximation that can be achieved with the resolvent model using deterministic expansion coefficients, no matter how their amplitudes and phases are chosen. Specifically, we will see that convergent approximations of the power spectral and cross-spectral densities are not possible when the statistical nature of the expansion coefficients is not respected.

A few specific choices of b_j are worth a closer look. First, we consider the case where the b_j values are chosen to have the correct magnitude, so that their power spectral densities $S_{b_j b_j} = |b_j|^2$ match the correct values $S_{\beta_j \beta_j}$. This could be accomplished, for example, using the approach of Gómez *et al.* (2016a). Equation (5.18) shows that the cross-spectral and power spectral densities depend also on off-diagonal terms $S_{\beta_j \beta_k}$. These will be replaced by the values $S_{b_j b_k} = b_j b_k^*$, and there is no reason to expect this to provide a good approximation. For example, in the case of uncorrelated expansion coefficients, the off-diagonal terms of $S_{\beta_j \beta_k}$ are zero, but $b_j b_k^*$ is non-zero unless either mode j or mode k has zero magnitude. Thus, even though the power spectral densities of the expansion coefficients are correct, the power spectral densities of the output will be incorrect. To make the error explicit, (5.25) can be manipulated into the form

$$\begin{aligned} \mathbf{s}_{yy}(\mathbf{x}, \mathbf{x}', f) &= \sum_{j=1}^{\infty} \hat{\mathbf{u}}_j(\mathbf{x}, f) \hat{\mathbf{u}}_j^*(\mathbf{x}', f) \sigma_j^2(f) |b_j(f)|^2 \\ &+ \sum_{j=1}^{\infty} \sum_{\substack{k=1 \\ k \neq j}}^{\infty} \hat{\mathbf{u}}_j(\mathbf{x}, f) \hat{\mathbf{u}}_k^*(\mathbf{x}', f) \sigma_j(f) \sigma_k(f) b_j(f) b_k^*(f). \end{aligned} \quad (5.26)$$

Using the power-spectral-density-based expansion coefficients, the first term in (5.26) is correct, while the second term is incorrect no matter how many terms are included in the expansion.

An alternative approach would be to choose the b_j coefficients such that the resolvent-mode expansion optimally matches the data, as defined by the inner product $\langle \cdot, \cdot \rangle_y$, in the limit of high N_r . Spectral POD theory tells us that the optimal representation of the data under the rank-one limitation is given by the leading SPOD mode. Therefore, the optimal b_j values are those that reconstruct the leading SPOD mode at each frequency. These values can be obtained by directly projecting the leading SPOD mode onto the resolvent modes, giving

$$b_j(f) = \frac{\sqrt{\lambda_1(f)}}{\sigma_j(f)} \langle \boldsymbol{\psi}_1, \hat{\mathbf{u}}_j \rangle_y. \quad (5.27)$$

By construction, these expansion coefficients guarantee that the first SPOD mode at each frequency is recovered when enough terms are retained in the resolvent-mode expansion. Even though this is the optimal expansion possible using a deterministic

vector of expansion coefficients, it will provide a good approximation of the flow only if $\lambda_1(f) \gg \lambda_2(f)$ over the range of relevant frequencies. Therefore, the quality of the resolvent-mode expansion using optimal (or any other) deterministic coefficients is dependent first and foremost on the low-rank nature of the cross-spectral density tensor rather than of the resolvent operator.

In contrast, convergent approximations of the flow statistics can be achieved if the statistical nature of the resolvent-mode expansion coefficients is incorporated at the outset. Accordingly, we propose a practical method for approximating the required cross-spectral densities $S_{\beta_j\beta_k}$ using the leading SPOD modes. The resolvent-mode expansion coefficients are related to SPOD modes by comparing the SPOD and resolvent expansions of the cross-spectral density tensor given in (5.20). By taking two successive inner products of both expansions with respect to $\hat{\mathbf{u}}_j$ and $\hat{\mathbf{u}}_k$ and dividing by $\sigma_j(f)$ and $\sigma_k(f)$, we obtain the expression

$$S_{\beta_j\beta_k} = \sum_{n=1}^{\infty} \frac{\lambda_n(f)}{\sigma_j(f)\sigma_k(f)} \gamma_{nj}(f) \gamma_{nk}^*(f), \quad (5.28)$$

where

$$\gamma_{jk}(f) = \langle \boldsymbol{\psi}_j, \hat{\mathbf{u}}_k \rangle_y \quad (5.29)$$

is the projection between the j th SPOD mode and the k th resolvent output mode.

Using the first N_s terms in (5.28) to approximate $S_{\beta_j\beta_k}$ and a resolvent-mode expansion with N_r terms leads to the approximation of the second-order flow statistics,

$$\mathbf{S}_{yy}(\mathbf{x}, \mathbf{x}', f) \approx \sum_{n=1}^{N_s} \lambda_n(f) \left(\sum_{j=1}^{N_r} \gamma_{nj}(f) \hat{\mathbf{u}}_j(\mathbf{x}, f) \right) \left(\sum_{k=1}^{N_r} \gamma_{nk}(f) \hat{\mathbf{u}}_k(\mathbf{x}', f) \right)^*. \quad (5.30)$$

From the definition of γ_{jk} and the fact that the input resolvent modes are complete, we have that

$$\sum_{j=1}^{N_r} \gamma_{nj}(f) \hat{\mathbf{u}}_j(\mathbf{x}, f) \rightarrow \boldsymbol{\psi}_n(\mathbf{x}, f) \quad \text{as } N_r \rightarrow \infty. \quad (5.31)$$

Therefore, using the first N_s SPOD modes to approximate $S_{\beta_j\beta_k}$ leads to a resolvent-mode expansion capable of capturing the first N_s SPOD modes, if enough resolvent modes are retained. The approximation is therefore convergent as N_s and N_r are increased, and if the SPOD spectrum falls off rapidly, accurate approximations of the flow are possible using only a few SPOD modes to compute the expansion coefficients. If a sufficiently high value of N_s is used, then the quality of the approximation that can be obtained with a limited number of resolvent modes is determined by the convergence of the limit in (5.31). It is notable that the resolvent gains do not appear explicitly in (5.30) or (5.31). Instead, the efficacy of the resolvent expansion is determined by the efficiency of the leading resolvent modes in providing a basis for the leading SPOD modes.

6. Examples

This section contains two example problems that demonstrate our theoretical results as well as the overall utility of SPOD.

6.1. Complex Ginzburg–Landau equation

Our first example consists of the linearized complex Ginzburg–Landau equation, which is frequently used as a model for instabilities in spatially evolving flows. The equation can be written in the form of (5.3) with

$$\mathcal{A} = -\nu \frac{\partial}{\partial x} + \gamma \frac{\partial^2}{\partial x^2} + \mu(x). \quad (6.1)$$

The spatial dependence of the solution is a consequence of the parameter $\mu(x)$, for which we adopt the quadratic form

$$\mu(x) = (\mu_0 - c_\mu^2) + \frac{\mu_2}{2}x^2 \quad (6.2)$$

used previously by several authors (Hunt & Crighton 1991; Bagheri *et al.* 2009; Chen & Rowley 2011). We take $\mu_0 = 0.23$, and all other parameters in (6.1) and (6.2) are set to the values used by Bagheri *et al.* (2009). The resulting model is globally stable (all eigenvalues of \mathcal{A} have negative real part) but is susceptible to non-modal growth due to the non-normality of \mathcal{A} , which within the context of resolvent analysis leads to gains much larger than one. The input and output operators and inner-product matrices (\mathcal{B} , \mathcal{C} , \mathbf{W}_y and \mathbf{W}_η respectively) are all set to unity. The value of μ_0 was chosen so that the gain of the leading resolvent mode at its peak frequency is 100 times larger than the gain of the second mode, which is a typical value for real flows.

The equations are discretized with a pseudo-spectral approach using Hermite polynomials, as in Bagheri *et al.* (2009) and Chen & Rowley (2011). The collocation points correspond to the roots of the first N Hermite polynomials; following Bagheri *et al.* (2009), we use $N = 220$. This leads to a computational domain $x \in [-85, 85]$, which is large enough to mimic an infinite domain. Since the collocation points are unevenly distributed, the discretized form of the inner product would contain matrices different from the identity. To avoid this, we define the input η and output y on a different uniformly spaced grid. This is accomplished by setting the discretized forms of the input and output matrices B and C so that they represent an interpolation from the uniform grid to the Hermite grid and from the Hermite grid to the uniform grid respectively. On the uniform input/output grid, the discrete inner-product matrices \mathbf{W}_y and \mathbf{W}_η are the identity matrix.

To generate data for our analysis, the discretized equations are stochastically excited in the time domain using forcing terms with prescribed statistics. In the following two subsections, we consider cases of white-noise and correlated forcing respectively.

6.1.1. White-noise forcing

We begin by forcing the linearized Ginzburg–Landau equations with band-limited white noise that is spatially uncorrelated. This forcing is realized by setting the value at each grid point and at each discrete time instance to a random complex number with uniformly distributed phase (between 0 and 2π) and normally distributed amplitude with unit variance. This signal is then low-pass filtered in time using a 10th-order finite-impulse-response filter with a cutoff frequency equal to 60% of the Nyquist frequency. Additionally, the forcing is spatially limited to the interior portion of the domain using an exponential envelope of the form $\exp[(x/L)^p]$ with $L = 60$ and $p = 10$.

The equations are integrated using a fourth-order embedded Runge–Kutta method (Shampine & Reichelt 1997), and a total of 40 000 snapshots of the solution are




FIGURE 2. (Colour online) Comparison of the first four resolvent gains σ_j^2 (lines) and SPOD eigenvalues λ_j (symbols) as a function of frequency for the Ginzburg–Landau model forced with white noise: (○, blue) mode 1; (□, green) mode 2; (△, red) mode 3; (▽, cyan) mode 4. For each mode, every other SPOD eigenvalue has been omitted from the plot to improve readability.

collected with spacing $\Delta t = 0.5$, leading to a Nyquist frequency of $f_{Nyquist} = 1$. A large number of snapshots are required for this problem because of the slow statistical convergence of the white-noise forcing; problems with correlated forcing, as is typical in real flows, tend to require fewer snapshots. Spectra and SPOD modes are computed as outlined in § 3 with a Hann window function, 75% overlap and $N_f = 384$.

Since the forcing is white and the input inner product has unit weight, the conditions are met under which resolvent and SPOD modes are theoretically equivalent. The first four SPOD eigenvalues and resolvent gains are shown in figure 2 as a function of frequency. For consistency with previous publications, frequencies are reported in terms of the angular frequency $\omega = 2\pi f$. Overall, the SPOD eigenvalues agree with the resolvent gains, but two minor differences can be observed which are related to the finite level of convergence of the spectral estimates used to approximate the SPOD modes. First, the SPOD eigenvalues are not completely smooth as a function of frequency. This is indicative of the remaining uncertainty in the spectral estimates, which can be further reduced by averaging over more data ensembles. Second, there is a small overshoot in the value of the second SPOD eigenvalue at the peak frequency of the first mode. This overshoot is related to spectral leakage, which can be reduced by increasing the frequency resolution of the spectral estimates by increasing the length of the data ensembles. When estimating the cross-spectral density using a data set of finite length, increasing the number of ensembles reduces the length of each ensemble and *vice versa*. Therefore, decreasing one of the two types of errors tends to increase the other. The proper compromise will in general be problem-dependent. As a general rule, SPOD modes are more difficult to converge for flows that exhibit sharp spectral peaks.




FIGURE 3. Weighted mode shapes in ω – x space for the Ginzburg–Landau equation forced with white noise. (a,d,g) Resolvent modes $\sigma_j(\omega)|\hat{\mathbf{u}}_j(x, \omega)|$. (b,e,h) Spectral POD modes $\sqrt{\lambda_j(\omega)}|\boldsymbol{\psi}_j(x, \omega)|$. (c,f,i) Space-only POD modes $|\hat{\mathbf{a}}_j(\omega)| |\boldsymbol{\phi}_j(x)|$.

The structure of the first three resolvent and SPOD modes is depicted in figure 3 as a function of x and ω (space-only POD results are also shown and will be discussed later). Specifically, we plot the weighted quantities $\sigma_j(\omega)|\hat{\mathbf{u}}_j(x, \omega)|$ and $\sqrt{\lambda_j(\omega)}|\boldsymbol{\psi}_j(x, \omega)|$ to clearly show where each mode is active in ω – x space. The contour levels in each plot are distributed between zero and the maximum value of the weighted mode.

The first resolvent and SPOD modes are nearly indistinguishable. The suboptimal SPOD modes likewise closely mimic their corresponding resolvent modes, although the limited smoothness as a function of frequency is again evident. The match between the resolvent and SPOD modes can be quantified using the projection γ_{jk} , defined in (5.29). This is shown for the peak frequency of the leading resolvent mode, $\omega = -0.6$, in figure 4(a). The diagonal ($j = k$) entries dominate up to $j = k = 12$, indicating that the first 12 resolvent and SPOD modes are nearly the same. Finally, figure 4(b) shows the expansion-coefficient cross-spectral density tensor $\mathcal{S}_{\beta_j\beta_k}$ for the same frequency, normalized by its maximum value. As expected, it is nearly diagonal, which is the root cause of the approximate equivalence of the resolvent and SPOD modes for this problem.

We also briefly compare the SPOD results with those obtained using space-only POD. Figure 3(c,f,i) shows the ω – x distribution of the first three POD modes. The frequency distribution was obtained by computing the power spectral densities of the time-dependent expansion coefficients $a_j(t)$ using Welch's method, which is similar to the method of Cammilleri *et al.* (2013) except that we use proper spectral averaging rather than a single DFT to obtain frequency information. The quantity plotted is




FIGURE 4. Spectral POD/resolvent-mode projection coefficient (a) and resolvent-mode expansion-coefficient cross-spectral density (b) at the frequency of highest gain for the Ginzburg–Landau model forced with white noise. The first 12 SPOD and resolvent modes are nearly identical (diagonal γ_{jk}) because the resolvent-mode expansion coefficients are almost completely uncorrelated (diagonal $S_{\beta_j \beta_k}$).

then $|\hat{a}_j(\omega)| |\phi_j(x)|$. As expected, each space-only POD mode represents a range of frequencies. The first space-only POD mode is relatively similar to the leading SPOD mode, but subsequent POD modes represent structures that are quite different. Two things stand out. First, since the spatial shape of each space-only POD mode does not depend on frequency, they are unable to capture the frequency-dependent shape of the resolvent modes that is captured by the SPOD modes. Second, the frequency-dependent magnitude of each suboptimal space-only POD mode exhibits a clear dip at the frequencies at which previous modes peak. Overall, the SPOD modes provide a better representation than the space-only POD modes of the underlying dynamics and non-normal growth mechanisms, which are completely described by the resolvent modes for this problem since the forcing is white.

Finally, we demonstrate the result from § 2.3 that space-only POD modes do not represent structures that are uncorrelated from one another in time. Figure 5 shows the correlation $C_{a_1 a_j}^{POD}(\tau) = E\{a_1(t) a_j^*(t + \tau)\}$ of the first POD expansion coefficient with the first three expansion coefficients ($j = 1, 2, 3$) as a function of the temporal lag τ , scaled by $\sqrt{\lambda_1 \lambda_j}$. The solid and broken lines show the real part and magnitude of the correlations respectively. At zero time lag, the scaled autocorrelation $C_{a_1 a_1}^{POD}(0)/\lambda_1$ is one and the cross-correlations $C_{a_1 a_k}^{POD}(0)$ are zero. Both of these values follow theoretically from (2.9). However, for $\tau > 0$, the autocorrelation quickly drops and the cross-correlations become non-zero. Accordingly, the space-only POD modes do not represent structures that evolve coherently in time.

6.1.2. Correlated forcing

Next, we force the linearized Ginzburg–Landau equations with forcing terms that are spatially correlated. This forcing is generated by convolving the same band-limited white-noise signal used previously with a kernel of the form

$$g(x, x') = \frac{1}{\sqrt{2\pi}\sigma_\eta} \exp \left[-\frac{1}{2} \left(\frac{x - x'}{\sigma_\eta} \right)^2 \right] \exp \left[i2\pi \frac{x - x'}{\lambda_\eta} \right], \quad (6.3)$$

where σ_η is the standard deviation of the envelope and λ_η is the wavelength of the filter. This leads to a forcing that is white-in-time up to the cutoff frequency but that




FIGURE 5. (Colour online) Temporal correlation of the first three space-only POD expansion coefficients with the first coefficient (see (2.24)) for the Ginzburg–Landau model forced with white noise. Real part of the correlation for (—, blue) mode 1; (---, green) mode 2; (-·-·-, red) mode 3. The light dotted lines show the magnitude of each correlation. Since the autocorrelation of the first mode decays and the cross-correlation with the suboptimal modes becomes non-zero for $\tau > 0$, the space-only POD modes do not represent structures that evolve coherently in space and time.

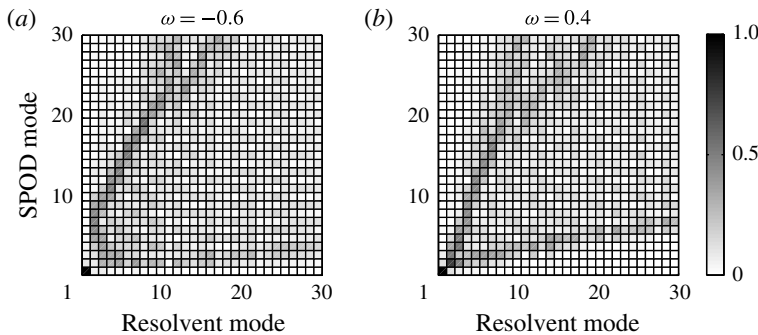


FIGURE 6. Spectral POD/resolvent projection coefficients γ_{jk} at two frequencies for the Ginzburg–Landau model forced by spatially correlated input. Since the tensors are not diagonal, the resolvent and SPOD modes are not the same in this case.

has non-zero spatial correlation in the form of (6.3) but with σ_η replaced with $\sqrt{2}\sigma_\eta$. This form of the forcing statistics is qualitatively similar to those of the nonlinear forcing terms in real flows, such as a turbulent jet (Towne *et al.* 2017). We use $\sigma_\eta = 4$ and $\lambda_\eta = 20$. The statistics of the response converge more rapidly in this case due to the correlated forcing, so we use a smaller number of snapshots, $N = 10\,000$, than in the uncorrelated case.

The correlation of the forcing leads to differences between the output resolvent and SPOD modes. This is clearly demonstrated by the γ_{jk} projection coefficients, which are shown for $\omega = -0.6$ and 0.4 in figure 6. The first of these frequencies is near the peak frequency of the leading resolvent mode where $\sigma_1^2 \approx 560$ and $\sigma_1^2/\sigma_2^2 \approx 100$, while the second is away from the peak and $\sigma_1^2 \approx 20$ and $\sigma_1^2/\sigma_2^2 \approx 10$. In both cases, the diagonal character of γ_{jk} observed for the white-noise forcing (figure 4a) is no longer observed, indicating that the SPOD and resolvent modes are no longer the same.

As discussed in §5, this can be attributed to the fact that the correlated forcing produces correlated resolvent-mode expansion coefficients. The magnitude of the cross-spectral density of the expansion coefficients $S_{\beta_j \beta_k}$ is shown in figure 7(a,d) for the

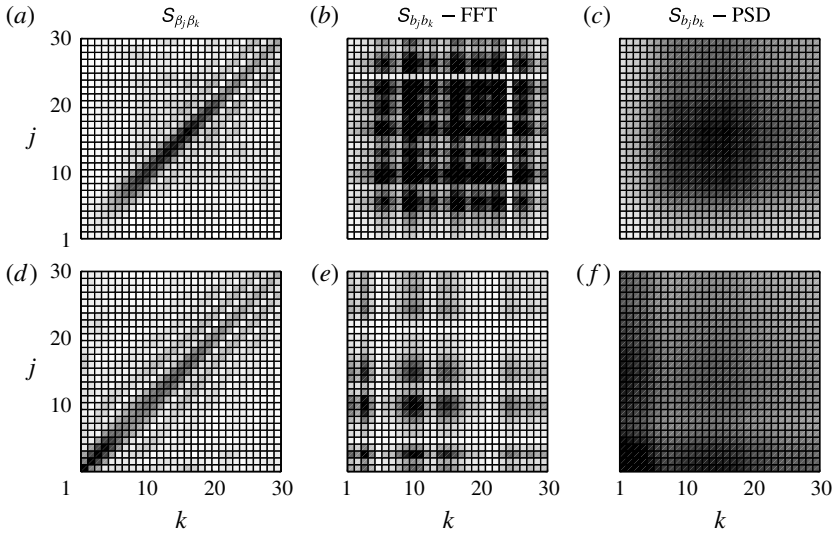


FIGURE 7. Cross-spectral density of the resolvent-mode expansion coefficients for (a,b,c) $\omega = -0.6$ and (d,e,f) $\omega = 0.4$. (a,d) Correct values. (b,e) Values implied by deterministic coefficients based on a single long-time DFT of the solution. (c,f) Values implied by deterministic coefficients with correct power spectral density.

same two frequencies. The strictly diagonal correlation tensor observed for the white-noise forcing (figure 4b) has been replaced by a more complicated banded structure, indicating substantial correlations between different expansion coefficients.

The expansion-coefficient correlations are not captured by standard approaches to computing the expansion coefficients that treat them as deterministic quantities. For example, figures 7(b,e) and 7(c,f) show the correlations implied using the DFT- and power spectral density (PSD)-based methods discussed in § 5.4, both of which lead to large errors away from the main correlation bands.

As indicated by (5.18), these errors in the expansion-coefficient cross-spectral density lead to errors in the resolvent-mode reconstruction of the output statistics. Figure 8 compares the true power spectral density of the output as a function of ω and x with the power spectral densities obtained from resolvent-mode reconstructions of the flow with $N_r = 1, 5, 10$ and 30 modes and with the expansion coefficients specified in four different ways. Figure 8(a) shows the true power spectral density, which does not depend on N_r . Figure 8(b) shows the power spectral density obtained using the correct values of $S_{\beta_j \beta_k}$, computed using (5.28). Figures 8(c) and 8(d) show the power spectral densities obtained using the DFT- and PSD-based deterministic expansion coefficients respectively, and figure 8(e) shows results for the optimal b_j values given by (5.27). The contour levels in all of the panels are the same and span six orders of magnitude, with the upper bound set to the maximum value of the true power spectral density.

All four resolvent-mode reconstructions produce similar power spectral densities for $N_r = 1$, but the DFT-based expansion is noticeably noisy due to the uncertainty implicit in this approach. As the number of resolvent modes included in the expansion increases, the power spectral density obtained using the proper $S_{\beta_j \beta_k}$ values converges to the true power spectral density. In contrast, none of the three methods using deterministic expansion coefficients converge to the true PSD due to their incorrect implied correlations. For the DFT- and PSD-based methods, the reconstructed power

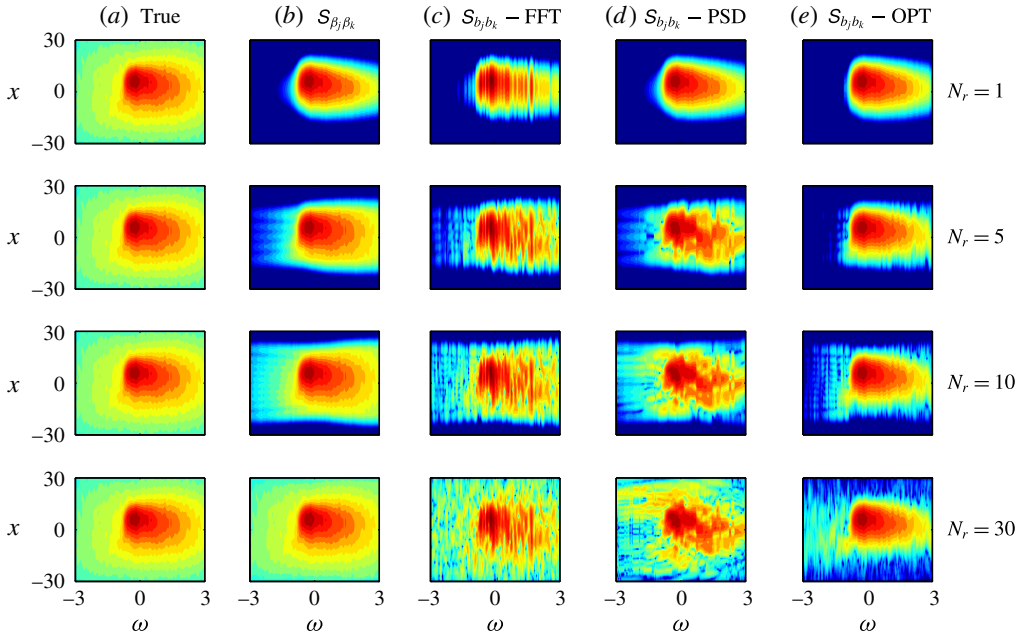


FIGURE 8. (Colour online) Resolvent-mode reconstructions of the PSD of the solution of the Ginzburg–Landau equation for spatially correlated forcing using different expansion coefficients. Rows: the number of resolvent modes in the expansion increases from top to bottom. Columns: (a) True PSD (same in every row). (b) Reconstruction with correct statistical expansion coefficients. (c) Reconstruction with deterministic expansion coefficients based on a single long-time DFT. (d) Reconstruction with deterministic PSD-based expansion coefficients. (e) Reconstruction with optimal deterministic expansion coefficients.

spectral densities improve with increasing N_r in low-energy regions of ω – x space but do not improve in high-energy regions. In fact, the reconstruction of the high-energy regions actually worsens with increasing N_r for the PSD-based expansion coefficients due to accumulation of the error terms from (5.26) as more modes are included in the expansion. On the other hand, the inclusion of more terms in the resolvent-mode expansion using the optimal b_j values has less effect on the low-energy regions but does not degrade the approximation of the high-energy regions. This is because the resolvent-mode reconstruction using these expansion coefficients is converging towards the leading SPOD mode at each frequency, and these modes effectively represent high-energy regions by construction. The expansion with optimal b_j values is well converged using 30 resolvent modes, so this is the best approximation of the second-order output statistics that can be obtained using deterministic expansion coefficients.

6.2. Turbulent jet

Our second example is a subsonic turbulent jet issued from a round convergent-straight nozzle. The Mach number, temperature ratio and Reynolds number of the jet are $M = U_j/c_j = 0.4$, $T_j/T_\infty = 1$ and $Re = \rho_j U_j D / \mu_j \approx 10^6$ respectively, where U is the velocity, c is the speed of sound, T is the temperature, ρ is the density, D is the nozzle diameter, μ is the dynamic viscosity, and the subscripts j and ∞ indicate mean conditions at the nozzle exit and in the ambient far field respectively.




FIGURE 9. (Colour online) Instantaneous snapshot of the Mach 0.4 turbulent jet. Greyscale: pressure fluctuations. Colour: vorticity magnitude.

We use data from a large-eddy simulation (LES) as input for the various empirical modal decomposition methods discussed in this paper. The simulation was performed using the compressible flow solver ‘Charles’ developed at Cascade Technologies (Brès *et al.* 2017a), which solves the spatially filtered compressible Navier–Stokes equations on unstructured grids using a finite-volume method and third-order Runge–Kutta time integration. The simulation used approximately sixteen million control volumes and was run for a duration of 2000 acoustic time units (tc_∞/D). The numerical methods, geometry and grid are identical to those used in a previous simulation at a higher Mach number, which was validated using double-blind comparisons with an extensive set of experimental measurements (Brès *et al.* 2015, 2016, 2017b).

The available LES database consists of 10 000 snapshots of the jet (velocities, density and pressure) sampled every 0.2 acoustic time units on a structured cylindrical output grid that approximately mirrors the underlying LES resolution and extends a distance of 30 jet diameters in the streamwise direction and six jet diameters in the radial direction. The azimuthal vorticity magnitude and pressure fluctuations from one of the snapshots are shown in figure 9. These snapshots are used to compute SPOD, space-only POD, DMD and DFT modes. Since the azimuthal coordinate is periodic, the snapshots can be decomposed into azimuthal Fourier modes, and each azimuthal mode can be treated independently within the context of the modal decompositions. For the sake of brevity, we exclusively present results for the axisymmetric component, and all modes are visualized using the pressure field.

6.2.1. Spectral POD modes

First, we study the SPOD modes of the jet. Spectral POD has been applied to jets before using data from both experiments (Glauser *et al.* 1987; Arndt *et al.* 1997; Citriniti & George 2000; Gudmundsson & Colonius 2011; Sinha *et al.* 2014; Semeraro *et al.* 2016a) and simulations (Towne *et al.* 2015; Schmidt *et al.* 2017b). These studies have used a variety of different definitions of the quantity of interest \mathbf{q} , spatial domain Ω and inner-product weight \mathbf{W} . Here, we use the full state vector $\mathbf{q} = [\rho, u_x, u_r, u_\theta, p]$ in the domain described previously, and we define the inner-product weight such that the induced norm is equivalent to the compressible energy norm proposed by Chu (1965). The SPOD modes are estimated using the procedure outlined in § 3 using blocks containing $N_f = 256$ snapshots, leading to $N_b = 78$ realizations of the jet. A standard Hann window is used to reduce spectral leakage. Frequencies are reported in terms of the Strouhal number $St = fD/U_j$.

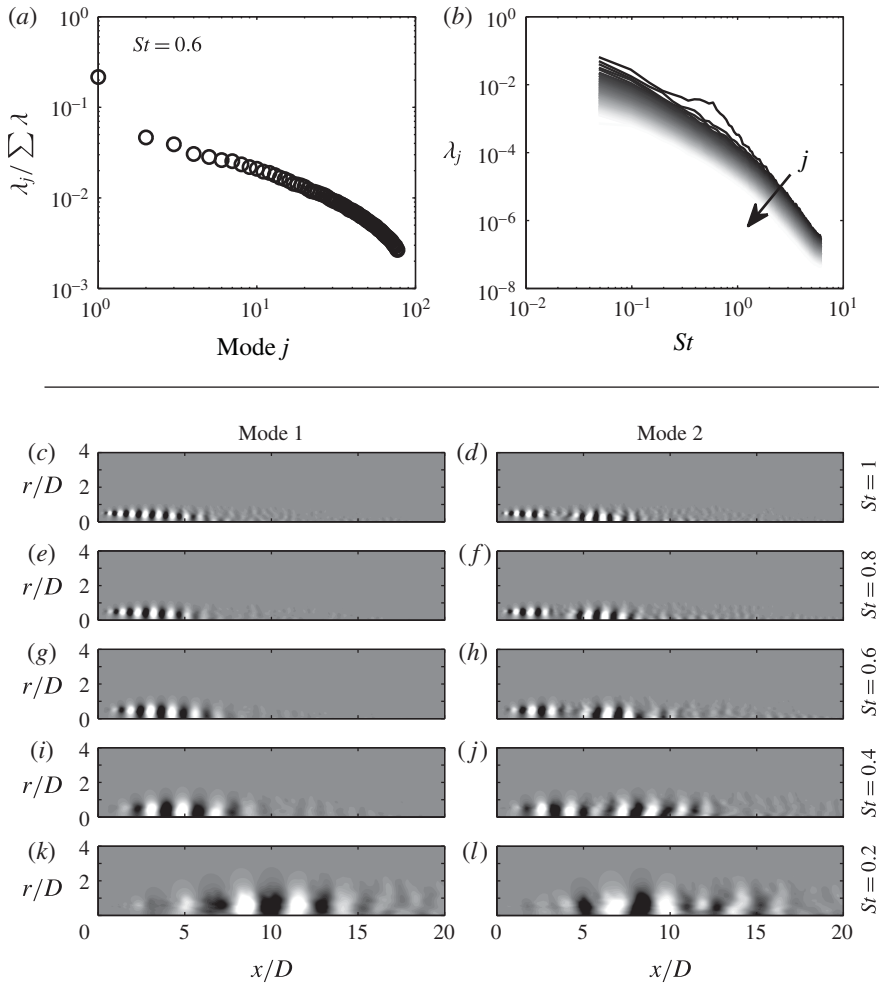


FIGURE 10. Spectral POD modes of the Mach 0.4 turbulent jet: (a) SPOD eigenvalues at $St = 0.6$, normalized by the total energy at that frequency; (b) SPOD eigenvalues as a function of frequency, normalized by the total flow energy; (c–l) real part of the pressure field of the first (c,e,g,i,k) and second (d,f,h,j,l) SPOD modes at the five indicated frequencies.

The SPOD modes are summarized in figure 10. The eigenvalues are depicted in two different ways in figure 10(a,b). Figure 10(a) shows the eigenvalues as a function of the mode number for $St = 0.6$. The eigenvalues are normalized by the total energy at this frequency, so each scaled eigenvalue represents the fraction of the energy at $St = 0.6$ described by that mode. The leading mode is substantially more energetic than the suboptimal modes and captures $\sim 22\%$ of the flow energy at this frequency. This is indicative of low-rank dynamics at this frequency (Schmidt *et al.* 2017b). The full spectrum of SPOD eigenvalues is shown in figure 10(b) as a function of the frequency. All eigenvalues have been normalized by the total flow energy integrated over all frequencies, so each one can be interpreted as the fraction of the total flow energy described by that mode. The shading of the curves varies linearly from black to white as the mode number increases from 1 to $N_b = 78$. The low-rank behaviour

observed at $St = 0.6$, indicated by a large gap between the first and second SPOD eigenvalues, persists in the range $0.3 \lesssim St \lesssim 1$.

The pressure fields of the first two SPOD modes are plotted for several frequencies in figure 10(c–l). All of the modes take the form of coherent wavepacket structures. The wavelength and spatial support of the wavepackets are strongly dependent on the frequency; both quantities increase with decreasing frequency. The wavepacket described by the first mode at each frequency has a single streamwise maximum. These wavepackets have been linked to the Kelvin–Helmholtz instability of the annular jet shear layer for $St \gtrsim 0.3$ (Suzuki & Colonius 2006; Gudmundsson & Colonius 2011) and to a superposition of non-normal modes for $St \lesssim 0.3$ (Jordan *et al.* 2017; Schmidt *et al.* 2017b). This link between the real turbulent jet and these simple concepts from linear stability theory provides a starting point for the construction of non-empirical reduced-order models. The wavepackets described by the second mode at each frequency have two streamwise maxima, and subsequent modes have increasingly complex structure. These modes optimally describe the variability in the shape of the wavepackets observed in the jet at different times (see § 6.2.3). These variations are central to the acoustic characteristics of the jet (Cavaliere *et al.* 2011; Cavaliere & Agarwal 2014).

Overall, the SPOD modes provide valuable insights that enhance physical understanding of the jet dynamics and motivate reduced-order models. In particular, the wavepackets represented by the SPOD modes have been shown to play a key role in generating acoustic radiation and provide a rigorous starting point for modelling and mitigating jet noise.

6.2.2. Space-only POD modes

Second, we examine the space-only POD modes of the jet. We will see that, in contrast to SPOD, the space-only POD modes provide little insight into the jet dynamics. The space-only POD eigenvalues are shown in figure 11(a). The eigenvalues decay slowly and give no indication of the low-rank behaviour of the jet revealed by SPOD. The power spectral density of each expansion coefficient is shown in figure 11(b). The contour levels are logarithmically distributed over three orders of magnitude, with the upper limit equal to the maximum value. The first approximately 10 modes are dominated by low frequencies that are below the relevant range for jet-noise research; even though these are the highest-energy structures in the jet, they are of little interest. All of the modes, and especially for $j \gtrsim 10$, contain contributions from a wide range of frequencies – each mode represents behaviour at many different time scales.

Because of this mixing of different time scales, the space-only POD modes cannot reproduce the orderly wavepacket structures identified by SPOD. The pressure fields of several modes are shown in figure 11(c–h). Modes 1 and 10 are dominated by large structures associated with the low frequencies that these modes represent. The higher modes have a disorganized appearance caused by the superposition of many different spatial scales associated with the wide range of time scales contained in each mode. Fragments of different wavepackets can be observed in these modes, but a clean separation of distinct structures is not achieved; as indicated by (2.28), the space-only POD modes are made up of a combination of many SPOD modes at different frequencies. This scale mixing makes the space-only POD modes far less helpful for understanding and modelling the jet dynamics. This is perhaps unsurprising in light of




FIGURE 11. Space-only POD modes of the Mach 0.4 turbulent jet: (a) space-only POD eigenvalues, normalized by the total flow energy; (b) PSD of the space-only POD expansion coefficients; (c–j) real part of the pressure field of space-only POD modes 1, 10, 20, 40, 60, 80, 100 and 200.

the analysis in § 2.3 showing that space-only POD modes do not represent physically meaningful coherent structures.

6.2.3. DMD modes

Third, we compare four different types of DMD modes. In all cases, we use the formulation given by (4.1) with different choices of the matrices \mathbf{X} and \mathbf{Y} . The first set of DMD modes is obtained using all 10 000 snapshots of the jet to define \mathbf{X} and \mathbf{Y} as in (4.2) without subtracting the mean – this is the standard way in which DMD would be applied to the jet database. The second set of DMD modes is obtained again using all 10 000 snapshots in the same way, but the mean is subtracted – this is equivalent to taking a single long-time DFT of the database. The third set of DMD modes is obtained from the ensemble DMD problem defined by (4.4) and the mean is subtracted from each flow realization – the DMD modes are thus the short-time DFT modes of each realization, as shown in § 4. The details of the ensemble definition are the same as described in § 6.2.1 for the SPOD computation. The fourth and final set of DMD modes is the SPOD modes constructed from the ensemble of short-time DFT modes.

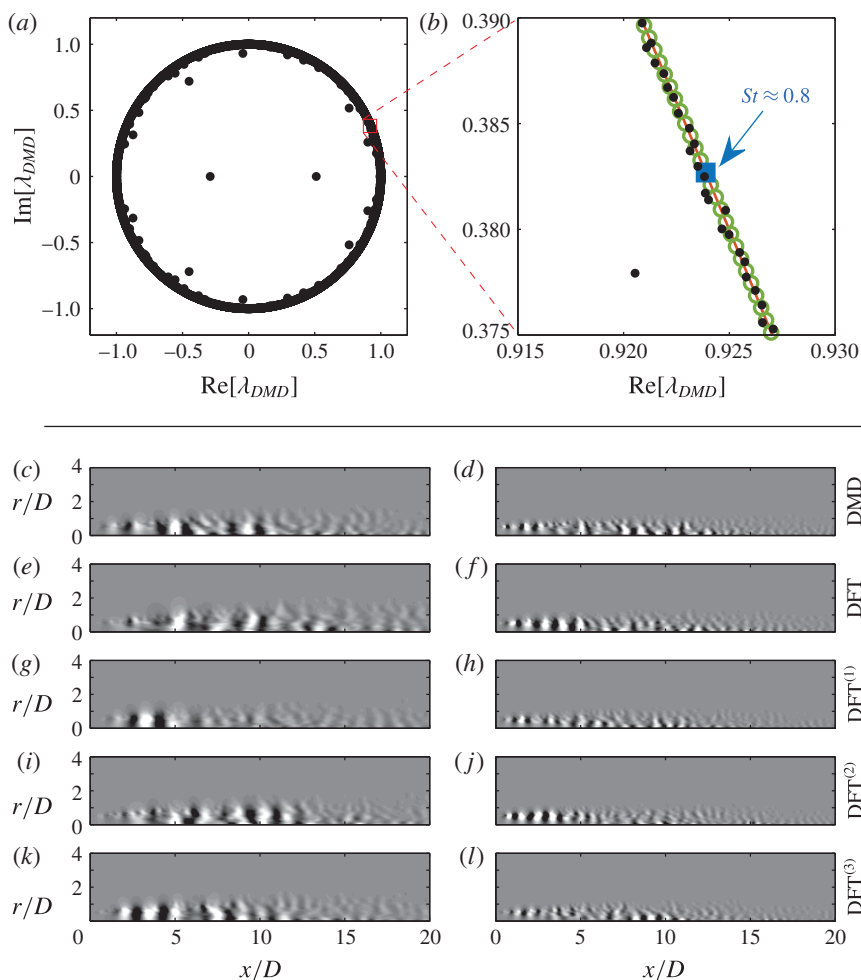


FIGURE 12. (Colour online) Dynamic mode decomposition/DFT modes of the Mach 0.4 turbulent jet: (a) eigenvalues for DMD without mean subtraction; (b) eigenvalues in the vicinity of $St = 0.8$ for (●) DMD without mean subtraction; (○, green) DMD with mean subtraction, i.e. a long-time DFT; (■, blue) ensemble DMD, i.e. short-time DFT; (c–l) real part of the pressure field of the different types of DMD modes; (c,d) DMD without mean subtraction; (e,f) long-time DFT; (g–l) three realizations of the short-time DFT. Panels (c,e,g,i,k) and (d,f,h,j,l) show the modes nearest to $St = 0.4$ and 0.8 respectively.

The DMD eigenvalues associated with the different sets of DMD modes are depicted in figure 12(a,b). Figure 12(a) shows the full set of eigenvalues for the standard DMD calculation in which the mean is not removed. With the exception of a small number of outliers, the eigenvalues are tightly clustered along the unit circle and therefore have nearly zero growth/decay rates. This is expected since the jet is a stationary flow. Figure 12(b) focuses on the portion of the unit circle corresponding to $St \approx 0.8$ and confirms that the growth/decay rates of the standard DMD modes are indeed small. Also shown in this plot are the DMD eigenvalues associated with the long-time and short-time DFT (only one short-time eigenvalue is visible). Most of the eigenvalues from the standard DMD calculation lie near one of the eigenvalues

from the long-time DFT, which shows that the standard DMD eigenvalues are nearly evenly spaced and suggests that each standard DMD mode is probably similar to one of the long-time DFT modes, even without mean subtraction.

The pressure fields of some of the DMD modes are shown in figure 12(c–l). Specifically, we show the modes from each different type of DMD that are nearest to the frequencies $St = 0.4$ and 0.8 . The standard DMD and long-time DFT modes are shown in figures 12(c,d) and 12(e,f) respectively. We see that these modes are quite similar, which is expected both from examination of the associated DMD eigenvalues and from the theoretical connection between DMD modes, Koopman modes and Fourier modes discussed in § 4. Three of the short-time DFT modes associated with the ensemble DMD problem are shown for the two frequencies in figure 12(g–l). Significant variability is observed between the different realizations, which is a consequence of the physical flow variability in different short-time intervals and the uncertainty inherent to the DFT. The SPOD modes shown in figure 10 provide an optimal basis for describing the ensemble of short-time DFT modes at each frequency. It is thus not surprising that traces of the wavepackets described by the first two SPOD modes can be observed in the short-time DFT modes and indeed in all of the different types of DMD modes. However, none of these other DMD/DFT modes clearly isolate the coherent wavepacket structures. Instead, they each show one of the many ways in which the coherent structures represented by the SPOD modes can coexist in one particular realization of the turbulent flow. Clearly, the elementary coherent structures provided by SPOD are more useful for understanding and modelling the jet than these random realizations provided by the other forms of DMD.

6.2.4. Resolvent modes

Fourth, we compute resolvent modes of the jet and compare them with the SPOD modes. Several authors have computed resolvent modes for jets in recent years (Garnaud *et al.* 2013; Jeun *et al.* 2016; Semeraro *et al.* 2016a,b; Schmidt *et al.* 2017b), and their properties have been shown to predict certain properties of the jet. Our focus here is on making comparisons with SPOD, as suggested by our analysis in § 5.

The jet is governed by the compressible Navier–Stokes equations, so the flow operator \mathcal{A} in (5.3) corresponds to these equations linearized about the mean flow computed from the LES data. The output operator \mathcal{C} selects the domain $\Omega = \{x, r \in [0, 30] \times [0, 6]\}$ in order to match the support of the LES snapshots used for the empirical methods. The input operator \mathcal{B} multiplies each component of η by the turbulent kinetic energy of the jet (again computed from the LES data). This choice is motivated by the analysis of Towne *et al.* (2017) and emphasizes the forcing terms in regions of the jet where turbulent fluctuations are of high enough amplitude to interact nonlinearly. The same compressible energy inner product (Chu 1965) as used for the POD modes is adopted for both the input and output spaces.

The equations are discretized using fourth-order finite differences on a grid with 950×250 points in x and r respectively. The computational domain includes the physical domain described above plus a surrounding sponge region that prevents waves from being reflected back into the computational domain at its boundaries. The overall numerical scheme is the same as in Schmidt *et al.* (2017a); additional details are available in that reference. As indicated by the above definition of Ω , the sponge regions carry zero weight in the discretized inner products. The numerical approximations of the leading resolvent output modes $\tilde{\mathbf{U}}$ and singular values Σ are

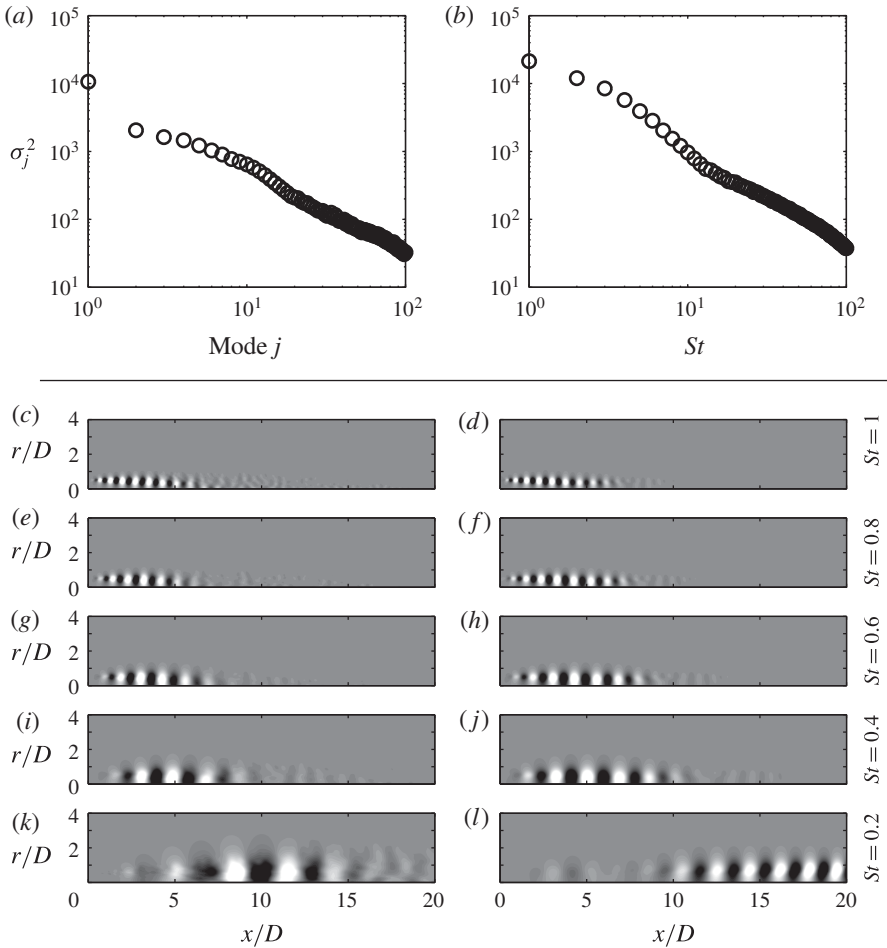


FIGURE 13. Resolvent modes of the Mach 0.4 turbulent jet: gain as a function of mode number for (a) $St = 0.6$ and (b) $St = 0.2$; (c–l) comparison between the leading SPOD mode (c,e,g,i,k) and the resolvent output mode (d,f,h,j,l) at the five indicated frequencies. The real part of the pressure field of each mode is shown.

calculated from the eigenvalue decomposition $\tilde{\mathbf{R}}\tilde{\mathbf{R}}^* = \tilde{\mathbf{U}}\Sigma^2\tilde{\mathbf{U}}^*$ using a standard Arnoldi method (see appendix C for a description of the preceding notation for the discretized equations).

The theory developed in § 5 provides a rigorous basis for comparing SPOD and resolvent modes. Figure 13(a,b) shows the gains of the first 100 resolvent modes for $St = 0.6$ and 0.2 . A large gain separation between the first and second modes is observed for $St = 0.6$ but not for $St = 0.2$. This mirrors the behaviour of the SPOD eigenvalues observed in figure 10(b). The leading resolvent modes provide an accurate approximation of the leading SPOD modes at frequencies that exhibit this low-rank behaviour. For example, the pressure fields of the leading SPOD and resolvent output modes are compared for several frequencies in figure 13(c–l). The low-rank behaviour is observed in the resolvent gain and SPOD eigenvalue spectra for all but the lowest frequency shown in the figure, and for these frequencies the resolvent modes approximate the corresponding SPOD modes with striking accuracy.

The only discernible difference is that the resolvent-mode wavepackets tend to be slightly longer. On the other hand, the leading resolvent mode at $St = 0.2$ is quite different from the corresponding SPOD mode. Similarly, the suboptimal resolvent modes at all frequencies (not shown) do not have a one-to-one correspondence with the suboptimal SPOD modes.

The theoretical connection between SPOD and resolvent modes enables us to glean useful information from these observations. The close match between the leading SPOD and resolvent modes for $St \gtrsim 0.3$ implies that the wavepackets described by these modes are robust and insensitive to correlations within the nonlinear forcing field. Conversely, the mismatch of lower-frequency and suboptimal modes indicates that it is necessary to account for correlated forcing in order to model the low-frequency jet dynamics as well as the wavepacket variability described by the suboptimal SPOD modes at all frequencies. Leveraging these ideas to better understand the jet dynamics and construct accurate reduced-order models is the subject of ongoing research (Schmidt *et al.* 2017b; Towne *et al.* 2017).

7. Conclusions

This paper explores a space–time formulation of POD for stationary flows called SPOD and its relationship with three other modal decomposition techniques. The main results of our analysis can be summarized by the following four statements about SPOD.

First, SPOD modes represent physically meaningful coherent structures in the sense that each mode evolves coherently in space and time – the part of the flow described by a particular SPOD mode is perfectly correlated with the part of the flow described by that same mode at all times and entirely uncorrelated with the part of the flow described by all other modes at all times. In contrast, modes obtained from a more common form of POD, which we call space-only POD, do not necessarily evolve coherently in time – the part of the flow described by a particular space-only POD mode is not necessarily correlated with the part of the flow described by the same mode at a later time, nor is it necessarily uncorrelated with the part of the flow described by a different mode at a later time. This essential difference can be traced back to the definition of the stochastic ensembles that form the starting point of each method; SPOD uses time-dependent flow realizations and thus retains dynamical information, whereas space-only POD discards all such information by treating each time instance of the flow as an independent realization of a random process.

Second, SPOD modes are optimally averaged DMD modes obtained from an ensemble DMD problem for stationary flows. The ensemble DMD problem is defined using the ‘exact DMD’ framework of Tu *et al.* (2014), and our result holds under the assumption of linearly consistent and zero-mean data ensembles. The first of these requirements corresponds to the condition under which DMD modes are expected to approximate Koopman modes, while the second assumption is justified by the zero growth rate of Koopman modes for stationary flows. We show that any linear combination of the DFT modes of the data ensembles are eigenvectors of this ensemble DMD problem. Spectral POD modes can be written as linear combinations of these DFT modes at each frequency, and are thus DMD modes. Specifically, SPOD modes are optimally averaged DMD modes that provide the best description of the statistical variability of the flow.

Third, SPOD modes are identical to resolvent modes in the special case in which the resolvent-mode expansion coefficients are uncorrelated. This result is based on

a statistical perspective on the resolvent-mode reconstruction of turbulent flows and on the observation that the resolvent-mode expansion coefficients are inherently statistical quantities for stationary turbulent flows. Uncorrelated expansion coefficients are usually associated with white-noise forcing of the linear dynamics. While the nonlinear perturbation terms of the Navier–Stokes equations are unlikely to be truly white for real flows, this has proven to be a reasonable first approximation for modelling a variety of flows. Accordingly, resolvent modes obtained from these models can be understood as non-empirical approximations of SPOD modes under the assumption of white-noise nonlinear forcing.

Fourth, SPOD modes provide fundamental insight into how the resolvent-mode expansion coefficients should be chosen so that the expansion properly captures the flow statistics. Specifically, we show the importance of properly accounting for the statistical nature of the correlated expansion coefficients and provide a method of computing these statistics using SPOD modes that leads to convergent approximations of the flow. All previous methods for obtaining the expansion coefficients assumed them to be deterministic quantities that can be entirely described by their amplitude and phase. We show that this results in a rank-one approximation of the second-order flow statistics, no matter how many resolvent modes are retained in the expansion. Since the optimal rank-one approximation of the cross-spectral density tensor is given by the leading SPOD mode at each frequency, the quality of the approximation that can be achieved using resolvent modes with deterministic expansion coefficients is governed by the low-rank nature of the second-order flow statistics rather than the resolvent operator.

These results are demonstrated using two example problems. The first is the linearized Ginzburg–Landau equation, which provides a simple model of a convectively unstable flow susceptible to non-modal growth. When forced with white noise, the SPOD modes computed from data and resolvent modes are nearly the same, as expected. Space-only POD modes computed from the same data do not capture these underlying dynamics and do not represent structures that evolve coherently in time. The Ginzburg–Landau equation is then forced with spatially coherent input, and the results are used to demonstrate the inability of resolvent-mode expansions with deterministic expansion coefficients to reconstruct the power spectral density of the solution. In contrast, these statistics converge when proper statistical expansion coefficients are employed.

The second example problem is a Mach 0.4 turbulent jet. The SPOD modes, computed from an LES database, isolate different space–time scales of the jet and provide insight into the low-rank dynamics of large-scale coherent wavepacket structures and how they might be modelled. In contrast, space-only POD modes obscure these low-rank wavepacket dynamics by jumbling together many different space–time scales in each mode. The DMD modes (without mean subtraction) and long-time DFT modes (i.e. DMD modes with mean subtraction) are similar, while the short-time DFT modes that are used to compute SPOD modes reveal the variability present in the jet. Each of these types of DMD modes represents one possible way in which the coherent wavepacket structures identified by SPOD can coexist in one of the many possible realizations of the turbulent jet, which highlights the advantage of SPOD for turbulent flows.

Overall, our results show that SPOD combines the advantages of space-only POD and DMD for stationary flows. Each SPOD mode represents a structure that is dynamic in the same sense as DMD modes, but that also accounts for and optimally describes the statistical variability of the turbulent flow. Consequently, SPOD modes

could provide improved robustness over space-only POD and DMD modes for modelling, estimation and control of turbulent flows.

Spectral POD also automatically solves two challenges in DMD that have spurred a number of variations of the method (Rowley & Dawson 2017). First, DMD modes have no inherent ranking, making it challenging to select a reduced set of modes for model reduction. Spectral POD, in contrast, gives a ranked set of modes in each frequency bin defined by the DFT rather than a collection of modes at many slightly different frequencies over the same frequency range. Figure 12(b) provides a clear illustration of this; SPOD replaces the assortment of DMD/DFT modes with a ranked set of modes at the frequency indicated by the blue square. Second, DMD can be sensitive to noisy data. Spectral POD automatically reduces the effect of noise as a consequence of the spectral estimation process, and convergence can always be improved by including more flow realizations.

The main disadvantage of SPOD is that it typically requires more data than either space-only POD or DMD. The spatial statistics required for space-only POD usually converge more rapidly than the space–time statistics required for SPOD. Standard applications of DMD use a single realization of a flow; SPOD requires at least a handful of realizations that are typically obtained from a longer time series. Additionally, if the mean is not subtracted, the frequencies identified by DMD are not constrained to be evenly spaced as they are for SPOD. This can be advantageous for laminar or transitional flows with sharp spectral peaks, since resolving peaks using evenly spaced points requires longer data series. Arbabi & Mezić (2017) recently proposed an alternative method for computing Koopman modes for flows with sharp spectral peaks via an iterative application of the DFT; a similar approach could be used to improve the performance of SPOD for these flows.

Acknowledgements

A.T. gratefully acknowledges support from NASA grant no. NNX15AU93A and from the NNSA Predictive Science Academic Alliance Program II, grant no. DE-NA0002373. O.T.S. acknowledges support from the German science foundation (DFG) grant no. 3114/1-1, and O.T.S. and T.C. acknowledge support from ONR-N0014-11-1-0753 and ONR-N00014-16-1-2445. The authors also thank Dr K. K. Chen for supplying portions of the code used for the Ginzburg–Landau problem and Dr G. A. Brès for providing the jet LES data. The LES study was supported by NAVAIR SBIR project, under the supervision of Dr J. T. Spyropoulos. The main LES calculations were carried out on CRAY XE6 machines at DoD HPC facilities in ERDC DSRC.

Appendix A. Derivation of the SPOD eigenvalue problem

The SPOD eigenvalue problem defined by (2.18) was first derived by Lumley (1967, 1970). Here, we offer an alternative derivation that we find to be more straightforward. Using (2.17), the correlation function $\mathbf{C}(\mathbf{x}, \mathbf{x}', t - t')$ can be written as

$$\mathbf{C}(\mathbf{x}, \mathbf{x}', t - t') = \int_{-\infty}^{\infty} \mathbf{S}(\mathbf{x}, \mathbf{x}', f) e^{i2\pi ft} e^{-i2\pi ft'} df. \quad (\text{A } 1)$$

Inserting this form into (2.13) leads to the following simplifications:

$$\int_{-\infty}^{\infty} \int_{\Omega} \int_{-\infty}^{\infty} \mathbf{S}(\mathbf{x}, \mathbf{x}', f) e^{i2\pi ft} e^{-i2\pi ft'} \mathbf{W} \boldsymbol{\phi}(\mathbf{x}', t') df d\mathbf{x}' dt' = \lambda \boldsymbol{\phi}(\mathbf{x}, t), \quad (\text{A } 2)$$

$$\int_{\Omega} \int_{-\infty}^{\infty} \mathbf{S}(\mathbf{x}, \mathbf{x}', f) \mathbf{W} \left[\int_{-\infty}^{\infty} \boldsymbol{\phi}(\mathbf{x}', t') e^{-i2\pi f t'} dt' \right] e^{i2\pi f t} df d\mathbf{x}' = \lambda \boldsymbol{\phi}(\mathbf{x}, t), \quad (\text{A } 3)$$

$$\int_{\Omega} \int_{-\infty}^{\infty} \mathbf{S}(\mathbf{x}, \mathbf{x}', f) \mathbf{W} \hat{\boldsymbol{\phi}}(\mathbf{x}', f) e^{i2\pi f t} df d\mathbf{x}' = \lambda \boldsymbol{\phi}(\mathbf{x}, t), \quad (\text{A } 4)$$

where $\hat{\boldsymbol{\phi}}(\mathbf{x}, f)$ is the temporal Fourier transform of $\boldsymbol{\phi}(\mathbf{x}, t)$.

For the solution ansatz $\boldsymbol{\phi}(\mathbf{x}, t) = \boldsymbol{\psi}(\mathbf{x}, f') e^{i2\pi f' t}$, we have $\hat{\boldsymbol{\phi}}(\mathbf{x}, f) = \boldsymbol{\psi}(\mathbf{x}, f') \delta(f - f')$. Substituting these into (A 4) gives

$$\int_{\Omega} \int_{-\infty}^{\infty} \mathbf{S}(\mathbf{x}, \mathbf{x}', f) \mathbf{W} \boldsymbol{\psi}(\mathbf{x}, f') \delta(f - f') e^{i2\pi f t} df d\mathbf{x}' = \lambda \boldsymbol{\psi}(\mathbf{x}, f') e^{i2\pi f' t} \quad (\text{A } 5)$$

or

$$\int_{\Omega} \mathbf{S}(\mathbf{x}, \mathbf{x}', f') \mathbf{W} \boldsymbol{\psi}(\mathbf{x}, f') e^{i2\pi f' t} d\mathbf{x}' = \lambda \boldsymbol{\psi}(\mathbf{x}, f') e^{i2\pi f' t}. \quad (\text{A } 6)$$

Multiplying both sides of (A 6) by $e^{-i2\pi f' t}$ gives the final SPOD eigenvalue problem (2.18).

Appendix B. Proof of ensemble DMD result

In this appendix, we prove that the modes of the ensemble DMD problem described in § 4.2 are the DFT modes of each realization for zero-mean linearly consistent data. We begin by writing (4.1) as

$$\mathbf{A}\mathbf{X} = \mathbf{Y}. \quad (\text{B } 1)$$

Tu *et al.* (2014) showed that (B 1) follows from (4.1) if and only if \mathbf{X} and \mathbf{Y} are linearly consistent. We therefore assume that the ensemble data matrices are linearly consistent and note that this is precisely the condition under which DMD modes are expected to approximate Koopman modes (Tu *et al.* 2014).

Since we wish to show that the eigenvalues of \mathbf{A} are the DFT modes of each flow realization, it is helpful to write \mathbf{X} and \mathbf{Y} in terms of these DFT modes. To do so, we use the matrix

$$\mathbf{F}_M \triangleq \frac{1}{\sqrt{M}} \begin{bmatrix} 1 & 1 & 1 & 1 & \dots & 1 \\ 1 & z & z^2 & z^3 & \dots & z^{M-1} \\ 1 & z^2 & z^4 & z^6 & \dots & z^{2(M-1)} \\ 1 & z^3 & z^6 & z^9 & \dots & z^{3(M-1)} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & z^{M-1} & z^{2(M-1)} & z^{3(M-1)} & \dots & z^{(M-1)(M-1)} \end{bmatrix}, \quad (\text{B } 2)$$

which when applied to a vector gives its DFT, i.e. $\hat{\mathbf{v}} = \mathbf{F}_M \mathbf{v}$ for $\mathbf{v} \in \mathbb{R}^M$. The scalar $z = e^{-i2\pi/M}$ is the primitive M th root of unity and the $1/\sqrt{M}$ scaling makes the DFT unitary and identical to (3.4) for rectangular window weights $w_j = 1$. Other windows, which are known to be beneficial in practice (Bendat & Piersol 2000), can be incorporated into the present analysis by using the windowed data to define each $\mathbf{Q}^{(n)}$. We wish to take the DFT of each row of the matrices $\mathbf{Q}^{(n)}$. To do so, we must apply \mathbf{F}_M to the conjugate transpose of each $\mathbf{Q}^{(n)}$ and apply a second conjugate transpose to the product to restore the original size of the matrix, $\hat{\mathbf{Q}}^{(n)} = (\mathbf{F}_M (\mathbf{Q}^{(n)})^*)^* = \mathbf{Q}^{(n)} \mathbf{F}_M^*$.

The first column of $\hat{\mathbf{Q}}^{(n)}$ gives the $k = 1$ term from (3.4), and so on. Since \mathbf{F}_M is unitary, it follows that $\mathbf{Q}^{(n)} = \hat{\mathbf{Q}}^{(n)} \mathbf{F}_M$. Inserting this into (4.4) gives

$$\mathbf{X} = [\hat{\mathbf{Q}}^{(1)} \mathbf{F}_M \mathbf{T}_X, \dots, \hat{\mathbf{Q}}^{(N_e)} \mathbf{F}_M \mathbf{T}_X] = [\hat{\mathbf{Q}}^{(1)} \hat{\mathbf{T}}_X, \dots, \hat{\mathbf{Q}}^{(N_e)} \hat{\mathbf{T}}_X], \quad (\text{B } 3a)$$

$$\mathbf{Y} = [\hat{\mathbf{Q}}^{(1)} \mathbf{F}_M \mathbf{T}_Y, \dots, \hat{\mathbf{Q}}^{(N_e)} \mathbf{F}_M \mathbf{T}_Y] = [\hat{\mathbf{Q}}^{(1)} \hat{\mathbf{T}}_Y, \dots, \hat{\mathbf{Q}}^{(N_e)} \hat{\mathbf{T}}_Y], \quad (\text{B } 3b)$$

where $\hat{\mathbf{T}}_{X,Y} = \mathbf{F}_M \mathbf{T}_{X,Y}$.

Since we have subtracted the mean from each flow realization, the zero-frequency ($k = 1$) DFT mode is identically zero, so the first column of $\hat{\mathbf{Q}}^{(n)}$ is zero for all n . Accordingly, we can remove the first column from each $\hat{\mathbf{Q}}^{(n)}$ and the first row from each instance of $\hat{\mathbf{T}}_X$ and $\hat{\mathbf{T}}_Y$ without changing \mathbf{X} and \mathbf{Y} . This is the key step that enables the remaining manipulations required to prove our result. We will denote the modified forms of $\hat{\mathbf{Q}}^{(n)}$, $\hat{\mathbf{T}}_X$ and $\hat{\mathbf{T}}_Y$ as $\tilde{\mathbf{Q}}^{(n)}$, $\tilde{\mathbf{T}}_X$ and $\tilde{\mathbf{T}}_Y$ respectively. It should be noted that while $\hat{\mathbf{T}}_X$ and $\hat{\mathbf{T}}_Y$ were rectangular matrices, $\tilde{\mathbf{T}}_X$ and $\tilde{\mathbf{T}}_Y$ are square $(M - 1) \times (M - 1)$ matrices. Specifically,

$$\tilde{\mathbf{T}}_X = \frac{1}{\sqrt{M}} \begin{bmatrix} 1 & z & z^2 & z^3 & \dots & z^{M-2} \\ 1 & z^2 & z^4 & z^6 & \dots & z^{2(M-2)} \\ 1 & z^3 & z^6 & z^9 & \dots & z^{3(M-2)} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & z^{M-1} & z^{2(M-1)} & z^{3(M-1)} & \dots & z^{(M-1)(M-2)} \end{bmatrix}, \quad (\text{B } 4a)$$

$$\tilde{\mathbf{T}}_Y = \frac{1}{\sqrt{M}} \begin{bmatrix} z & z^2 & z^3 & z^4 & \dots & z^{M-1} \\ z^2 & z^4 & z^6 & z^8 & \dots & z^{2(M-1)} \\ z^3 & z^6 & z^9 & z^{12} & \dots & z^{3(M-1)} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ z^{M-1} & z^{2(M-1)} & z^{3(M-1)} & z^{4(M-1)} & \dots & z^{(M-1)(M-1)} \end{bmatrix}. \quad (\text{B } 4b)$$

These matrices can be factored as

$$\tilde{\mathbf{T}}_X = \mathbf{F}_{M-1} \text{diag}([1, z, z^2, \dots, z^{M-2}]), \quad (\text{B } 5a)$$

$$\tilde{\mathbf{T}}_Y = \text{diag}([z, z^2, \dots, z^{M-1}]) \mathbf{F}_{M-1} \text{diag}([1, z, \dots, z^{M-2}]), \quad (\text{B } 5b)$$

which will prove useful in what follows.

Using these modified matrices, the input and output data matrices can be written as

$$\mathbf{X} = [\tilde{\mathbf{Q}}^{(1)} \tilde{\mathbf{T}}_X, \dots, \tilde{\mathbf{Q}}^{(N_e)} \tilde{\mathbf{T}}_X], \quad (\text{B } 6a)$$

$$\mathbf{Y} = [\tilde{\mathbf{Q}}^{(1)} \tilde{\mathbf{T}}_Y, \dots, \tilde{\mathbf{Q}}^{(N_e)} \tilde{\mathbf{T}}_Y]. \quad (\text{B } 6b)$$

Inserting (B 6) into (B 1) gives the expression

$$\mathbf{A}[\tilde{\mathbf{Q}}^{(1)}, \dots, \tilde{\mathbf{Q}}^{(N_e)}] \begin{bmatrix} \tilde{\mathbf{T}}_X & & \\ & \ddots & \\ & & \tilde{\mathbf{T}}_X \end{bmatrix} = [\tilde{\mathbf{Q}}^{(1)}, \dots, \tilde{\mathbf{Q}}^{(N_e)}] \begin{bmatrix} \tilde{\mathbf{T}}_Y & & \\ & \ddots & \\ & & \tilde{\mathbf{T}}_Y \end{bmatrix}. \quad (\text{B } 7)$$

Since \mathbf{F}_{M-1} is unitary and powers of z are non-zero, (B 5a) shows that $\tilde{\mathbf{T}}_X$ is the product of two invertible matrices and is therefore itself invertible. As a result, equation (B 7) can be written in the form given by (4.5) with $\Lambda_{\text{DMD}} \triangleq \tilde{\mathbf{T}}_Y \tilde{\mathbf{T}}_X^{-1}$. Finally, using (B 5), Λ_{DMD} reduces to the form given by (4.6).

Appendix C. Discretized SPOD/resolvent-mode relations

In practice, SPOD and resolvent modes are always approximated on discrete grids. To help make the results of § 5 as accessible and useful as possible, in this appendix, we write the discrete forms of some of the key equations relating these different types of modes. The discretized forms of all continuous variables and operators will be represented by the corresponding upright symbols, e.g. \mathbf{S}_{yy} becomes \mathbf{S}_{yy} . Time or frequency dependences can be inferred from the continuous variables and these arguments are suppressed for brevity.

We begin by writing the discretized form of the linearized flow equations and output as

$$\frac{d\chi'}{dt} - \mathbf{A}\chi' = \mathbf{B}\eta \quad (\text{C } 1)$$

and

$$\mathbf{y} = \mathbf{C}\chi' \quad (\text{C } 2)$$

respectively (to be clear, \mathbf{A} is the discretization of \mathcal{A} , not the DMD matrix). Note that the input and output do not need to be discretized on the same grid as the linearized flow operator and state vector, since the input and output matrices \mathbf{B} and \mathbf{C} can be used to interpolate from one grid to another. We will assume that the discretized output \mathbf{y} is defined on the same grid as the flow data used to estimate SPOD modes.

The inner products on the input and output spaces are approximated as $\langle \hat{\mathbf{y}}_1, \hat{\mathbf{y}}_2 \rangle_{\mathbf{y}} = \hat{\mathbf{y}}_2^* \mathbf{W}_{\mathbf{y}} \hat{\mathbf{y}}_1$ and $\langle \hat{\boldsymbol{\eta}}_1, \hat{\boldsymbol{\eta}}_2 \rangle_{\boldsymbol{\eta}} = \hat{\boldsymbol{\eta}}_2^* \mathbf{W}_{\boldsymbol{\eta}} \hat{\boldsymbol{\eta}}_1$ respectively. It is important to note that the positive-definite Hermitian matrices $\mathbf{W}_{\mathbf{y}}$ and $\mathbf{W}_{\boldsymbol{\eta}}$ must account for both the weight matrices and the numerical quadrature of the integrals in (5.10) and (5.11).

The discretized resolvent operator is

$$\mathbf{R} = \mathbf{C}(i2\pi f \mathbf{I} - \mathbf{A})^{-1} \mathbf{B}. \quad (\text{C } 3)$$

If $\mathbf{W}_{\mathbf{y}} = \mathbf{W}_{\boldsymbol{\eta}} = \mathbf{I}$, then the resolvent modes are given by the singular value decomposition

$$\mathbf{R} = \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^*. \quad (\text{C } 4)$$

The singular values appear within the diagonal positive-semidefinite matrix $\boldsymbol{\Sigma}$, and the input and output modes are contained in the columns of the orthonormal matrices \mathbf{V} and \mathbf{U} respectively. When either of the weight matrices $\mathbf{W}_{\mathbf{y}}$ or $\mathbf{W}_{\boldsymbol{\eta}}$ is not the identity, the resolvent modes can be recovered from the singular value decomposition of a weighted matrix,

$$\tilde{\mathbf{R}} = \mathbf{W}_{\mathbf{y}}^{1/2} \mathbf{R} \mathbf{W}_{\boldsymbol{\eta}}^{-1/2} = \tilde{\mathbf{U}} \boldsymbol{\Sigma} \tilde{\mathbf{V}}^*. \quad (\text{C } 5)$$

The singular values again appear on the diagonal of $\boldsymbol{\Sigma}$, and the input and output modes are contained in the columns of the matrices $\mathbf{V} = \mathbf{W}_{\boldsymbol{\eta}}^{-1/2} \tilde{\mathbf{V}}$ and $\mathbf{U} = \mathbf{W}_{\mathbf{y}}^{-1/2} \tilde{\mathbf{U}}$ respectively. The resolvent operator is recovered as

$$\mathbf{R} = \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^* \mathbf{W}_{\boldsymbol{\eta}}. \quad (\text{C } 6)$$

The resolvent-mode expansion of the Fourier-transformed output is then $\hat{\mathbf{y}} = \mathbf{U} \boldsymbol{\Sigma} \boldsymbol{\beta}$, where $\boldsymbol{\beta} = \mathbf{V}^* \mathbf{W}_{\boldsymbol{\eta}} \hat{\boldsymbol{\eta}}$. The discretized cross-spectral density tensor is thus

$$\mathbf{S}_{yy} = E\{\hat{\mathbf{y}} \hat{\mathbf{y}}^*\} = \mathbf{U} \boldsymbol{\Sigma} \mathbf{S}_{\boldsymbol{\beta}\boldsymbol{\beta}} \boldsymbol{\Sigma} \mathbf{U}^*, \quad (\text{C } 7)$$

where $\mathbf{S}_{\boldsymbol{\beta}\boldsymbol{\beta}} = E\{\boldsymbol{\beta} \boldsymbol{\beta}^*\} = \mathbf{V}^* \mathbf{W}_{\boldsymbol{\eta}} \mathbf{S}_{\boldsymbol{\eta}\boldsymbol{\eta}} \mathbf{W}_{\boldsymbol{\eta}} \mathbf{V}$ and $\mathbf{S}_{\boldsymbol{\eta}\boldsymbol{\eta}} = E\{\hat{\boldsymbol{\eta}} \hat{\boldsymbol{\eta}}^*\}$.

Using (C7) along with the data-based SPOD estimates discussed in §3, the output cross-spectral density can be written in terms of SPOD or resolvent modes as

$$\mathbf{S}_{yy} = \mathbf{\Psi} \mathbf{\Lambda} \mathbf{\Psi}^* = \mathbf{U} \mathbf{\Sigma} \mathbf{S}_{\beta\beta} \mathbf{\Sigma}^* \mathbf{U}^*. \quad (\text{C } 8)$$

This is the discrete equivalent of (5.18), which is the central expression in §5. If the resolvent-mode expansion coefficients are uncorrelated, then $\mathbf{S}_{\beta\beta} = \mathbf{I}$ and (C8) reduces to

$$\mathbf{S}_{yy} = \mathbf{\Psi} \mathbf{\Lambda} \mathbf{\Psi}^* = \mathbf{U} \mathbf{\Sigma}^2 \mathbf{U}^*, \quad (\text{C } 9)$$

and so $\mathbf{\Sigma}^2 = \mathbf{\Lambda}$ and $\mathbf{U} = \mathbf{\Psi}$. More generally, the expansion coefficients can be written in terms of SPOD modes as

$$\mathbf{S}_{\beta\beta} = \mathbf{\Sigma}^{-1} \mathbf{G} \mathbf{\Lambda} \mathbf{G}^* \mathbf{\Sigma}^{-1}, \quad (\text{C } 10)$$

where $\mathbf{G} = \mathbf{U}^* \mathbf{W}_y \mathbf{\Psi}$ is the discrete equivalent of γ_{ij} defined in (5.29). These equations are analogous to (5.28) and (5.29). Equivalently, $\mathbf{S}_{\beta\beta}$ can be conveniently written in terms of the data matrix $\hat{\mathbf{Q}}$ used to define the discrete SPOD eigenvalue problem. If we define $\mathbf{E} = \mathbf{\Sigma}^{-1} \mathbf{U}^* \mathbf{W}_y \hat{\mathbf{Q}}$, then $\mathbf{S}_{\beta\beta} = \mathbf{E} \mathbf{E}^*$.

If the statistical vector $\boldsymbol{\beta}$ is replaced by a deterministic vector \mathbf{b} , then the approximation of the output cross-spectral density matrix is $\mathbf{S}_{yy} = (\mathbf{U} \mathbf{\Sigma} \mathbf{b})(\mathbf{U} \mathbf{\Sigma} \mathbf{b})^*$, which is clearly a rank-one matrix. The optimal deterministic expansion coefficients, i.e. the ones that reconstruct the leading SPOD mode, are given by the vector $\mathbf{b}^{\text{opt}} = \mathbf{\Sigma}^{-1} \mathbf{G}_1 \lambda_1^{-1/2}$, where \mathbf{G}_1 is the first column of \mathbf{G} and λ_1 is the first entry of the diagonal matrix $\mathbf{\Lambda}$.

REFERENCES

- ABREU, L. I., CAVALIERI, A. V. G. & WOLF, W. R. 2017 Coherent hydrodynamic waves and trailing-edge noise. *AIAA Paper* 2017-3173.
- ARAYA, D. B., COLONIUS, T. & DABIRI, J. O. 2017 Transition to bluff-body dynamics in the wake of vertical-axis wind turbines. *J. Fluid Mech.* **813**, 346–381.
- ARBABI, H. & MEZIĆ, I. 2017 Study of dynamics in unsteady flows using Koopman mode decomposition. *Phys. Rev. Fluids* **2**, 124402.
- ARNDT, R. E. A., LONG, D. F. & GLAUSER, M. N. 1997 The proper orthogonal decomposition of pressure fluctuations surrounding a turbulent jet. *J. Fluid Mech.* **340**, 1–33.
- AUBRY, N. 1991 On the hidden beauty of the proper orthogonal decomposition. *Theoret. Comput. Fluid Dyn.* **2** (5), 339–352.
- AUBRY, N., HOLMES, P., LUMLEY, J. L. & STONE, E. 1988 The dynamics of coherent structures in the wall region of a turbulent boundary layer. *J. Fluid Mech.* **192**, 115–173.
- BAGHERI, S., HENNINGSON, D. S., HOEPFFNER, J. & SCHMID, P. J. 2009 Input–output analysis and control design applied to a linear model of spatially developing flows. *Appl. Mech. Rev.* **62** (2), 020803.
- BENDAT, J. S. & PERSOL, A. G. 2000 *Random Data: Analysis and Measurement Procedures*, 3rd edn. John Wiley & Sons.
- BENEDDINE, S., SIPP, D., ARNAULT, A., DANDOIS, J. & LESSHAFFT, L. 2016 Conditions for validity of mean flow stability analysis. *J. Fluid Mech.* **798**, 485–504.
- BERKOOZ, G., HOLMES, P. & LUMLEY, J. L. 1993 The proper orthogonal decomposition in the analysis of turbulent flows. *Annu. Rev. Fluid Mech.* **25** (1), 539–575.
- BRAUD, C., HEITZ, D., ARROYO, G., PERRET, L., DELVILLE, J. & B., J.-P. 2004 Low-dimensional analysis, using POD, for two mixing layer–wake interactions. *Intl J. Heat Fluid Flow* **25** (3), 351–363.

- BRÈS, G. A., HAM, F. E., NICHOLS, J. W. & LELE, S. K. 2017a Unstructured large-eddy simulations of supersonic jets. *AIAA J.* **55** (4), 1164–1184.
- BRÈS, G. A., JAUNET, J., LE RALLIC, M., JORDAN, P., COLONIUS, T. & LELE, S. K. 2015 Large eddy simulation for jet noise: the importance of getting the boundary layer right. *AIAA Paper* 2015-2535.
- BRÈS, G. A., JAUNET, V., LE RALLIC, M., JORDAN, P., TOWNE, A., SCHMIDT, O., COLONIUS, T., CAVALIERI, A. V. G. & LELE, S. K. 2016 Large eddy simulation for jet noise: azimuthal decomposition and intermittency of the radiated sound. *AIAA Paper* 2016-3050.
- BRÈS, G. A., JORDAN, P., JAUNET, V., LE RALLIC, M., CAVALIERI, A. V. G., TOWNE, A., LELE, S. K., COLONIUS, T. & SCHMIDT, O. T. 2017b Importance of the nozzle-exit boundary-layer state in subsonic turbulent jets. *J. Fluid Mech.* (in press).
- CAMMILLERI, A., GUÉNIAT, F., CARLIER, J., PASTUR, L., MÉMIN, E., LUSSEYRAN, F. & ARTANA, G. 2013 POD-spectral decomposition for fluid flow analysis and model reduction. *Theoret. Comput. Fluid Dyn.* **27** (6), 787–815.
- CAVALIERI, A. V. G. & AGARWAL, A. 2014 Coherence decay and its impact on sound radiation by wavepackets. *J. Fluid Mech.* **748**, 399–415.
- CAVALIERI, A. V. G., JORDAN, P., AGARWAL, A. & GERVAIS, Y. 2011 Jittering wave-packet models for subsonic jet noise. *J. Sound Vib.* **330** (18), 4474–4492.
- CHATTERJEE, A. 2000 An introduction to the proper orthogonal decomposition. *Curr. Sci.* **78** (7), 808–817.
- CHEN, K. K. & ROWLEY, C. W. 2011 H_2 optimal actuator and sensor placement in the linearised complex Ginzburg–Landau system. *J. Fluid Mech.* **681**, 241–260.
- CHEN, K. K., TU, J. H. & ROWLEY, C. W. 2012 Variants of dynamic mode decomposition: boundary condition, Koopman, and Fourier analyses. *J. Nonlinear Sci.* **22** (6), 887–915.
- CHEN, X. & KAREEM, A. 2005 Proper orthogonal decomposition-based modeling, analysis, and simulation of dynamic wind load effects on structures. *J. Engng Mech.* **131** (4), 325–339.
- CHU, B.-T. 1965 On the energy transfer to small disturbances in fluid flow (Part I). *Acta Mech.* **1** (3), 215–234.
- CITRINITI, J. H. & GEORGE, W. K. 2000 Reconstruction of the global velocity field in the axisymmetric mixing layer utilizing the proper orthogonal decomposition. *J. Fluid Mech.* **418**, 137–166.
- CORDIER, L. & BERGMANN, M. 2008 Proper orthogonal decomposition: an overview. In *Lecture Series 2002-04, 2003-03 and 2008-01 on Post-Processing of Experimental and Numerical Data*. Von Karman Institute for Fluid Dynamics.
- DELVILLE, J., UKEILEY, L., CORDIER, L., BONNET, J. P. & GLAUSER, M. 1999 Examination of large-scale structures in a turbulent plane mixing layer. Part 1. Proper orthogonal decomposition. *J. Fluid Mech.* **391**, 91–122.
- DERGHAM, G., SIPP, D. & ROBINET, J.-CH. 2013 Stochastic dynamics and model reduction of amplifier flows: the backward facing step flow. *J. Fluid Mech.* **719**, 406–430.
- FARRELL, B. F. & IOANNOU, P. J. 1993 Stochastic forcing of the linearized Navier–Stokes equations. *Phys. Fluids* **5** (11), 2600–2609.
- FARRELL, B. F. & IOANNOU, P. J. 1996 Generalized stability theory. Part I: autonomous operators. *J. Atmos. Sci.* **53** (14), 2025–2040.
- FARRELL, B. F. & IOANNOU, P. J. 2001 Accurate low-dimensional approximation of the linear dynamics of fluid flow. *J. Atmos. Sci.* **58** (18), 2771–2789.
- GARNAUD, X., LESSHAFFT, L., SCHMID, P. J. & HUERRE, P. 2013 The preferred mode of incompressible jets: linear frequency response analysis. *J. Fluid Mech.* **716**, 189–202.
- GEORGE, W. K. 1988 Insight into the dynamics of coherent structures from a proper orthogonal decomposition dy. In *International Seminar on Wall Turbulence*.
- GEORGE, W. K. 2017 A 50-year retrospective and the future. In *Whither Turbulence and Big Data in the 21st Century?* pp. 13–43. Springer.
- GEORGE, W. K., BEUTHER, P. D. & LUMLEY, J. L. 1978 Processing of random signals. In *Proceedings of the Dynamic Flow Conference 1978 on Dynamic Measurements in Unsteady Flows*, pp. 757–800. Springer.

- GLAUSER, M. N., LEIB, S. J. & GEORGE, W. K. 1987 *Coherent Structures in the Axisymmetric Turbulent Jet Mixing Layer*, pp. 134–145. Springer.
- GÓMEZ, F., BLACKBURN, H. M., RUDMAN, M., SHARMA, A. S. & MCKEON, B. J. 2016a On the coupling of direct numerical simulation and resolvent analysis. In *Progress in Turbulence VI*, pp. 87–91. Springer.
- GÓMEZ, F., BLACKBURN, H. M., RUDMAN, M., SHARMA, A. S. & MCKEON, B. J. 2016b A reduced-order model of three-dimensional unsteady flow in a cavity based on the resolvent operator. *J. Fluid Mech.* **798**, R2.
- GORDEYEV, S. V. & THOMAS, F. O. 2000 Coherent structure in the turbulent planar jet. Part 1. Extraction of proper orthogonal decomposition eigenmodes and their self-similarity. *J. Fluid Mech.* **414**, 145–194.
- GUDMUNDSSON, K. & COLONIUS, T. 2011 Instability wave models for the near-field fluctuations of turbulent jets. *J. Fluid Mech.* **689**, 97–128.
- HEINZEL, G., RÜDIGER, A. & SCHILLING, R. 2002 Spectrum and spectral density estimation by the discrete Fourier transform (DFT), including a comprehensive list of window functions and some new flat-top windows. https://holometer.fnal.gov/GH_FFT.pdf (unpublished).
- HELLSTRÖM, L. H. O. & SMITS, A. J. 2014 The energetic motions in turbulent pipe flow. *Phys. Fluids* **26** (12), 125102.
- 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** (3), 283–290.
- HOLMES, P., LUMLEY, J. L., BERKOOZ, G. & ROWLEY, C. W. 2012 *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, 2nd edn. Cambridge University Press.
- HOLMES, P. J., LUMLEY, J. L., BERKOOZ, G., MATTINGLY, J. C. & WITTENBERG, R. W. 1997 Low-dimensional models of coherent structures in turbulence. *Phys. Rep.* **287** (4), 337–384.
- HUNT, R. E. & CRIGHTON, D. G. 1991 Instability of flows in spatially developing media. *Proc. R. Soc. Lond. A* **435**, 109–128.
- JEUN, J., NICHOLS, J. W. & JOVANOVIĆ, M. R. 2016 Input–output analysis of high-speed axisymmetric isothermal jet noise. *Phys. Fluids* **28** (4), 047101.
- JORDAN, P., ZHANG, M., LEHNASCH, G. & CAVALIERI, A. V. G. 2017 Modal and non-modal linear wavepacket dynamics in turbulent jets. *AIAA Paper* 2017-3379.
- JOVANOVIĆ, M. & BAMIEH, B. 2001 Modeling flow statistics using the linearized Navier–Stokes equations. In *Decision and Control, 2001. Proceedings of the 40th IEEE Conference*, vol. 5, pp. 4944–4949. IEEE.
- JOVANOVIĆ, M. R. & BAMIEH, B. 2005 Componentwise energy amplification in channel flows. *J. Fluid Mech.* **534**, 145–183.
- LANDAHL, M. T. & MOLLO CHRISTENSEN, E. 1992 *Turbulence and Random Processes in Fluid Mechanics*. Cambridge University Press.
- LIANG, Y. C., LEE, H. P., LIM, S. P., LIN, W. Z., LEE, K. H. & WU, C. G. 2002 Proper orthogonal decomposition and its applications. Part I: theory. *J. Sound Vib.* **252** (3), 527–544.
- LUMLEY, J. L. 1967 The structure of inhomogeneous turbulent flows. In *Atmospheric Turbulence and Radio Propagation* (ed. A. M. Yaglom & V. I. Tatarski), pp. 166–178. Nauka.
- LUMLEY, J. L. 1970 *Stochastic Tools in Turbulence*. Academic Press.
- MCKEON, B. J. & SHARMA, A. S. 2010 A critical-layer framework for turbulent pipe flow. *J. Fluid Mech.* **658**, 336–382.
- MEZIĆ, I. 2005 Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dyn.* **41** (1), 309–325.
- MOARREF, R. & JOVANOVIĆ, M. R. 2012 Model-based design of transverse wall oscillations for turbulent drag reduction. *J. Fluid Mech.* **707**, 205–240.
- MOARREF, R., JOVANOVIĆ, M. R., TROPP, J. A., SHARMA, A. S. & MCKEON, B. J. 2014 A low-order decomposition of turbulent channel flow via resolvent analysis and convex optimization. *Phys. Fluids* **26** (5), 051701.
- MOIN, P. & MOSER, R. D. 1989 Characteristic-eddy decomposition of turbulence in a channel. *J. Fluid Mech.* **200**, 471–509.

- MULA, S. M. & TINNEY, C. E. 2014 Classical and snapshot forms of the POD technique applied to a helical vortex filament. *AIAA Paper* 2015-3257.
- NOACK, B. R., AFANASIEV, K., MORZYŃSKI, M., TADMOR, G. & THIELE, F. 2003 A hierarchy of low-dimensional models for the transient and post-transient cylinder wake. *J. Fluid Mech.* **497**, 335–363.
- NOACK, B. R., STANKIEWICZ, W., MORZYŃSKI, M. & SCHMID, P. J. 2016 Recursive dynamic mode decomposition of transient and post-transient wake flows. *J. Fluid Mech.* **809**, 843–872.
- PICARD, C. & DELVILLE, J. 2000 Pressure velocity coupling in a subsonic round jet. *Intl J. Heat Fluid Flow* **21** (3), 359–364.
- PINIER, J. T., AUSSEUR, J. M., GLAUSER, M. N. & HIGUCHI, H. 2007 Proportional closed-loop feedback control of flow separation. *AIAA J.* **45** (1), 181–190.
- POPE, S. B. 2000 *Turbulent Flows*. Cambridge University Press.
- ROWLEY, C. W. 2005 Model reduction for fluids, using balanced proper orthogonal decomposition. *Intl J. Bifurcation Chaos* **15** (03), 997–1013.
- ROWLEY, C. W., COLONIUS, T. & MURRAY, R. M. 2004 Model reduction for compressible flows using POD and Galerkin projection. *Physica D* **189** (1), 115–129.
- ROWLEY, C. W. & DAWSON, S. T. M. 2017 Model reduction for flow analysis and control. *Annu. Rev. Fluid Mech.* **49**, 387–417.
- ROWLEY, C. W., MEZIĆ, I., BAGHERI, S., SCHLATTER, P. & HENNINGSON, D. S. 2009 Spectral analysis of nonlinear flows. *J. Fluid Mech.* **641**, 115–127.
- SCHMID, P. J. 2010 Dynamic mode decomposition of numerical and experimental data. *J. Fluid Mech.* **656**, 5–28.
- SCHMID, P. J. & HENNINGSON, D. S. 2001 *Stability and Transition in Shear Flows*, vol. 142. Springer Science & Business Media.
- SCHMID, P. J. & SESTERHENN, J. 2008 Dynamic mode decomposition of numerical and experimental data. In *Bull. Am. Phys. Soc., 61st APS meeting*, p. 208. San Antonio.
- SCHMID, P. J., VIOLATO, D. & SCARANO, F. 2012 Decomposition of time-resolved tomographic PIV. *Exp. Fluids* **52** (6), 1567–1579.
- SCHMIDT, O. T. 2017 An efficient streaming algorithm for spectral proper orthogonal decomposition. [arXiv:1711.04199](https://arxiv.org/abs/1711.04199).
- SCHMIDT, O. T., TOWNE, A., COLONIUS, T., CAVALIERI, A. V. G., JORDAN, P. & BRÈS, G. A. 2017a Wavepackets and trapped acoustic modes in a turbulent jet: coherent structure eduction and global stability. *J. Fluid Mech.* **825**, 1153–1181.
- SCHMIDT, O. T., TOWNE, A., RIGAS, G., COLONIUS, T. & BRÈS, G. A. 2017b Spectral analysis for jet turbulence. [arXiv:1711.06296](https://arxiv.org/abs/1711.06296).
- SEMERARO, O., BELLANI, G. & LUNDELL, F. 2012 Analysis of time-resolved PIV measurements of a confined turbulent jet using POD and Koopman modes. *Exp. Fluids* **53** (5), 1203–1220.
- SEMERARO, O., JAUNET, V., JORDAN, P., CAVALIERI, A. V. G. & LESSHAFFT, L. 2016a Stochastic and harmonic optimal forcing in subsonic jets. *AIAA Paper* 2016-2935.
- SEMERARO, O., LESSHAFFT, L., JAUNET, V. & JORDAN, P. 2016b Modeling of coherent structures in a turbulent jet as global linear instability wavepackets: theory and experiment. *Intl J. Heat Fluid Flow* **62**, 24–32.
- SHAMPINE, L. F. & REICHEL, M. W. 1997 The Matlab ODE suite. *SIAM J. Sci. Comput.* **18** (1), 1–22.
- SHARMA, A. S. & MCKEON, B. J. 2013 On coherent structure in wall turbulence. *J. Fluid Mech.* **728**, 196–238.
- SHARMA, A. S., MEZIĆ, I. & MCKEON, B. J. 2016 Correspondence between Koopman mode decomposition, resolvent mode decomposition, and invariant solutions of the Navier–Stokes equations. *Phys. Rev. Fluids* **1** (3), 032402.
- SIEBER, M., PASCHEREIT, C. O. & OBERLEITHNER, K. 2016 Spectral proper orthogonal decomposition. *J. Fluid Mech.* **792**, 798–828.
- SINHA, A., RODRIGUEZ, D., BRÈS, G. A. & COLONIUS, T. 2014 Wavepacket models for supersonic jet noise. *J. Fluid Mech.* **742**, 71–95.

- SIPP, D., MARQUET, O., MELIGA, P. & BARBAGALLO, A. 2010 Dynamics and control of global instabilities in open-flows: a linearized approach. *Appl. Mech. Rev.* **63** (3), 030801.
- SIROVICH, L. 1987 Turbulence and the dynamics of coherent structures. I. Coherent structures. *Q. Appl. Maths* **45** (3), 561–571.
- SIROVICH, L. 1989 Chaotic dynamics of coherent structures. *Physica D* **37** (1), 126–145.
- SUZUKI, T. & COLONIUS, T. 2006 Instability waves in a subsonic round jet detected using a near-field phased microphone array. *J. Fluid Mech.* **565**, 197–226.
- TAIRA, K., BRUNTON, S. L., DAWSON, S., 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.
- TINNEY, C. E. & JORDAN, P. 2008 The near pressure field of co-axial subsonic jets. *J. Fluid Mech.* **611**, 175–204.
- TOWNE, A., BRÈS, G. A. & LELE, S. K. 2016 Toward a resolvent-based statistical jet-noise model. In *Annual Research Briefs*, Center for Turbulence Research, Stanford University.
- TOWNE, A., BRÈS, G. A. & LELE, S. K. 2017 A statistical jet-noise model based on the resolvent framework. *AIAA Paper* 2017-3406.
- TOWNE, A., COLONIUS, T., JORDAN, P., CAVALIERI, A. V. G. & BRÈS, G. A. 2015 Stochastic and nonlinear forcing of wavepackets in a Mach 0.9 jet. *AIAA Paper* 2015-2217.
- TREFETHEN, L., TREFETHEN, A., REDDY, S. & DRISCOLL, T. 1993 Hydrodynamic stability without eigenvalues. *Science* **261** (5121), 578–584.
- TU, J. H., ROWLEY, C. W., LUCHTENBURG, D. M., BRUNTON, S. L. & KUTZ, J. N. 2014 On dynamic mode decomposition: theory and applications. *J. Comput. Dyn.* **1** (2), 391–421.
- TUTKUN, M. & GEORGE, W. K. 2017 Lumley decomposition of turbulent boundary layer at high Reynolds numbers. *Phys. Fluids* **29** (2), 020707.
- TUTKUN, M., JOHANSSON, P. B. V. & GEORGE, W. K. 2008 Three-component vectorial proper orthogonal decomposition of axisymmetric wake behind a disk. *AIAA J.* **46** (5), 1118.
- WELCH, P. 1967 The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms. *IEEE Trans. Audio Electroacoust.* **15** (2), 70–73.
- ZARE, A., JOVANOVIĆ, M. R. & GEORGIU, T. T. 2017 Colour of turbulence. *J. Fluid Mech.* **812**, 636–680.