

Frequency–time analysis, low-rank reconstruction and denoising of turbulent flows using SPOD

Akhil Nekkanti¹ and Oliver T. Schmidt^{1,†}

¹Department of Mechanical and Aerospace Engineering, University of California San Diego, CA 92093, USA

(Received 10 November 2020; revised 13 April 2021; accepted 26 July 2021)

Four different applications of spectral proper orthogonal decomposition (SPOD) are demonstrated on large-eddy simulation data of a turbulent jet. These are: low-rank reconstruction, denoising, frequency–time analysis and prewhitening. We demonstrate SPOD-based flow-field reconstruction using direct inversion of the SPOD algorithm (frequency-domain approach) and propose an alternative approach based on projection of the time series data onto the modes (time-domain approach). We further present a SPOD-based denoising strategy that is based on hard thresholding of the SPOD eigenvalues. The proposed strategy achieves significant noise reduction while facilitating drastic data compression. In contrast to standard methods of frequency–time analysis such as wavelet transform, a proposed SPOD-based approach yields a spectrogram that characterises the temporal evolution of spatially coherent flow structures. A convolution-based strategy is proposed to compute the time-continuous expansion coefficients. When applied to the turbulent jet data, SPOD-based frequency–time analysis reveals that the intermittent occurrence of large-scale coherent structures is directly associated with high-energy events. This work suggests that the time-domain approach is preferable for low-rank reconstruction of individual snapshots, and the frequency-domain approach for denoising and frequency–time analysis.

Key words: low-dimensional models, intermittency, jets

1. Introduction

The curse of dimensionality (see e.g. Meneveau, Lund & Moin 1992) in the analysis of large turbulent flow data has led to the development of a number of modal decomposition

† Email address for correspondence: oschmidt@ucsd.edu

techniques (Holmes *et al.* 2012). The primary utilities of these techniques are to extract the essential flow features and to provide a low-dimensional representation of the data. Most of these techniques seek modes that lie in the span of the snapshots that constitute the time-resolved data, and adhere to certain mathematical properties that define the decomposition. Arguably the most widely used technique is proper orthogonal decomposition (POD), introduced by Lumley (1967, 1970). A specific version of POD, the computationally inexpensive method of snapshots (Sirovich 1987; Aubry 1991), decomposes the flow field into spatial modes and temporal coefficients. Its modes optimally represent the data in terms of its variance, or energy, and are coherent in space and at zero time lag. Another popular method is the dynamic mode decomposition (DMD; Schmid 2010), which is rooted in Koopman theory (Rowley *et al.* 2009) and assumes an evolution operator that maps the flow field from one snapshot to its next. The DMD modes are characterised by a single frequency and linear amplification rate. Refer to the reviews by Taira *et al.* (2017) and Rowley & Dawson (2017) for summaries of various modal techniques.

Spectral proper orthogonal decomposition (SPOD) is the frequency-domain variant of POD and computes modes as estimates of the eigenvectors of the cross-spectral density (CSD) matrix. At each frequency, SPOD yields a set of orthogonal modes, ranked by energy. The mathematical framework underlying SPOD was first outlined by Lumley (1967, 1970). Early implementations of SPOD include Glauser, Leib & George (1987), Glauser & George (1992), Delville (1994), Arndt, Long & Glauser (1997), Picard & Delville (2000), Citriniti & George (2000) and Gordeyev & Thomas (2000). For statistically stationary flows, Towne, Schmidt & Colonius (2018) have demonstrated that SPOD combines the advantages of POD, namely optimality and orthogonality, and DMD, namely temporal monochromaticity. In this work, we demonstrate how these properties can be leveraged for different applications.

In the past, SPOD has been used to analyse a number of turbulent flows, including jets (Arndt *et al.* 1997; Gamard *et al.* 2002; Gamard, Jung & George 2004; Gordeyev & Thomas 2000, 2002; Jung, Gamard & George 2004; Iqbal & Thomas 2007; Tinney, Glauser & Ukeiley 2008a; Tinney, Ukeiley & Glauser 2008b; Schmidt *et al.* 2018; Pickering *et al.* 2019; Nekkanti & Schmidt 2021), the wake behind a disk (Johansson, George & Woodward 2002; Johansson & George 2006a,b; Tutkun, Johansson & George 2008; Ghate, Towne & Lele 2020; Nidhan *et al.* 2020), pipe flows (Hellström & Smits 2014; Hellström, Ganapathisubramani & Smits 2015; Hellström, Marusic & Smits 2016; Hellström & Smits 2017) and channel flows (Muralidhar *et al.* 2019). Several studies have shown that a significant amount of energy is captured by the first few modes at each frequency. For jets and disk wakes, Glauser *et al.* (1987), Citriniti & George (2000), Jung *et al.* (2004), Johansson & George (2006b) and Tinney *et al.* (2008b) have shown that the leading mode and first three modes capture at least 40% and 80% of the total energy, respectively. The above studies have in common that SPOD modes and eigenvalues are interpreted directly as physical structures and energies. The applications shown in this work, however, require a full or partial reconstruction of the data in the time domain (often after a manipulation of the expansion coefficients in the frequency domain). Partial reconstructions of the flow field from SPOD were previously shown by Citriniti & George (2000), Tinney *et al.* (2008b) and, more recently, Ghate *et al.* (2020). We demonstrate low-rank reconstruction using two approaches. One is by inverting the SPOD algorithm, which was previously employed by Citriniti & George (2000) in a similar manner, for which we present an alternative means of computation based on convolution in the time domain. The other is by taking an oblique projection of the data on the SPOD modes. The advantages and disadvantages of both approaches for different applications are discussed.

The first of these applications is denoising. Most experimental flow-field data exhibit measurement noise that hampers physical analysis. The computation of spatial derivatives required for quantities such as the vorticity or strain rate, for example, leads to particularly large errors. Another difficulty is that physically relevant small-scale structures may be concealed by noise. The most common experimental technique for multi-dimensional flow-field measurement is particle image velocimetry (PIV). Common techniques to remove noise from PIV data include spatial filtering (Discetti, Natale & Astarita 2013), temporal filtering based on Fourier truncation, and POD-based techniques (Raiola, Discetti & Ianiro 2015; Brindise & Vlachos 2017). Spatial filters are typically based on Gaussian smoothing (Discetti *et al.* 2013). Different temporal filters such as median filters (Son & Kihm 2001), Hampel filters (Fore *et al.* 2005), Wiener filters (Vétel, Garon & Pelletier 2011) and band-pass filters (Sciacchitano & Scarano 2014) have been used for denoising PIV data. A comparison of different spatial and temporal filters is presented, for example, in Vétel *et al.* (2011). As an alternative to these standard techniques, POD reconstruction has been used as a means of denoising through mode truncation. Low-dimensional reconstructions from standard POD have been applied for this purpose to flow past a backward-facing step (Kostas, Soria & Chong 2005), arterial flows (Charonko *et al.* 2010; Brindise *et al.* 2017), turbulent wakes (Raiola *et al.* 2015) and vortex rings (Stewart & Vlachos 2012; Brindise & Vlachos 2017). In this contribution, we demonstrate the use of SPOD for denoising on surrogate data obtained by imposing high levels of additive Gaussian noise on simulation data. We demonstrate that SPOD-based denoising combines certain advantages of temporal filters and standard POD-based denoising.

Owing to their chaotic nature, turbulent flows are characterised by high levels of intermittency. A common tool for the analysis of intermittent behaviour is frequency–time analysis, that is, the representation of the frequency content of a time signal as a function of time. This representation is particularly well suited in identifying events such as short-time interval of high or low energy, and in identifying their wave characteristics, i.e. frequencies or wavenumbers. Frequency–time analysis can be performed using several different signal-processing tools such as wavelet transforms (WT) (Farge 1992), the short-time Fourier transform (STFT) (Cohen 1995), the S-transform (Stockwell, Mansinha & Lowe 1996), the Hilbert–Huang transform (Huang *et al.* 1998) and the Wigner–Ville distribution (Boashash 1988). We note that STFT and WT are arguably the most widely used techniques in fluid mechanics. Frequency–time diagrams obtained from these methods are generally referred to as spectrograms, or as scalograms for WT. The STFT performs Fourier transforms on consecutive short segments of a time signal. It has been used, for example, in the analysis of blood flows (Izatt *et al.* 1997; Zhang *et al.* 2003), magnetohydrodynamics (Bale *et al.* 2005), aerodynamics (Samimy *et al.* 2007) and physical oceanography (Brown *et al.* 1989). The WT is based on the convolution of the time signal with a compact waveform, the so-called mother wavelet, that is scaled to represent different frequencies. Typical applications of the WT are found in atmospheric science (Gu & Philander 1995), oceanography (Meyers, Kelly & O’Brien 1993) and, most importantly for the present work, turbulence research (Farge 1992). In the latter context, they have been used to extract coherent structures (Farge, Schneider & Kevlahan 1999; Farge, Pellegrino & Schneider 2001), and to analyse their intermittency (Camussi & Guj 1997; Onorato, Camussi & Iuso 2000; Camussi 2002). All methods mentioned above are signal-processing techniques that are applied to one-dimensional data. Here, we expand on the ideas of Schmidt, Colonius & Brés (2017a) and Towne & Liu (2019), and analyse the intermittency of the entire flow field. The underlying idea is that the global dynamics of the entire flow field can be described in terms of a limited set of statistically

prevalent, most energetic coherent flow structures. For statistically stationary flows, such structures are distilled by SPOD, and their temporal dynamics is described by the SPOD expansion coefficients. Since SPOD is a frequency-domain technique, this idea leverages the fact that each SPOD mode is associated with a single frequency. Based on the two reconstruction techniques mentioned above, we apply, analyse and compare two variants of SPOD-based frequency–time analysis. The frequency-domain approach relies on direct inversion of the SPOD algorithm and was previously demonstrated by Towne & Liu (2019). This approach, however, becomes computationally intractable even for moderately sized, two-dimensional data. We show that this problem can be avoided by the convolution-based approach introduced herein.

Prewhitening is a post-processing technique used for trend detection that is commonly used in the atmospheric and geophysical sciences. It was first proposed by Von Storch (1999). Prewhitening was used, for example, for the detection of trends in temperature and precipitation data (Zhang *et al.* 2001), rainfall (Lacombe, McCartney & Forkuor 2012), teleconnections (Rodionov 2006) and hydrological flows (Khaliq *et al.* 2009; Serinaldi & Kilsby 2015). Technically, prewhitening is achieved by a filtering operation that results in a flat power spectrum to remove serial correlations. We show two different SPOD-based ways to achieve this goal in the frequency domain.

The remainder of this paper is organised as follows. Section 2 describes the two techniques for SPOD-based flow-field reconstruction. In § 3, we demonstrate these techniques on the numerical data of a turbulent jet. The four different applications, SPOD-based low-dimensional reconstruction, denoising, frequency–time analysis and prewhitening, are demonstrated in § 3.1, § 3.2, § 3.3 and § 3.4, respectively. Section 4 summarises this work.

2. Methodology

2.1. Spectral proper orthogonal decomposition

In the following, we provide an outline of a specific procedure of computing SPOD based on Welch’s method (Welch 1967) and emphasise aspects that are important in the context of data reconstruction. Refer to the work of Towne *et al.* (2018) for details of the derivation and mathematical properties, and Schmidt & Colonius (2020) for a practical introduction to the method.

Given a fluctuating flow field $\mathbf{q}_i = \mathbf{q}(t_i)$, where $i = 1, \dots, n_t$, which is obtained by subtracting the temporal mean $\bar{\mathbf{q}}$ from each snapshot of the data, we start by constructing a snapshot matrix

$$\mathbf{Q} = [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_{n_t}]. \quad (2.1)$$

Note that multi-dimensional data are cast into the form of a vector \mathbf{q}_i of length n , corresponding to the number of variables times the number of grid points. The instantaneous energy of each time instant, or snapshot, is expressed in terms of a spatial inner product

$$\|\mathbf{q}\|_x^2 = \langle \mathbf{q}, \mathbf{q} \rangle_x = \int_{\Omega} \mathbf{q}^*(x, t) \mathbf{W}(x) \mathbf{q}(x, t) dx, \quad (2.2)$$

where \mathbf{W} is a positive-definite Hermitian matrix that accounts for the component-wise weights, Ω the spatial domain of interest and $(\cdot)^*$ denotes the complex conjugate. The common form of space-only POD is obtained as the eigendecomposition of $\mathbf{Q}\mathbf{Q}^*\mathbf{W}$ and yields modes that are optimal in terms of (2.2). The SPOD, however, specialises POD

for statistically stationary processes and seeks modes that are optimal in terms of the space–time inner product

$$\|q\|_{x,t}^2 = \langle q, q \rangle_{x,t} = \int_{-\infty}^{\infty} \int_{\Omega} q^*(x, t) W(x) q(x, t) dx dt. \quad (2.3)$$

For statistically stationary data, it is natural to proceed in the frequency domain and solve the POD eigenvalue problem for the Fourier transformed two-point space–time correlation matrix, that is, the CSD matrix. To estimate the CSD, the data are segmented into n_{blk} overlapping blocks with n_{fft} snapshots in each of them, as

$$Q^{(k)} = [q_1^{(k)}, q_2^{(k)}, \dots, q_{n_{fft}}^{(k)}]. \quad (2.4)$$

If the blocks overlap by n_{ovlp} snapshots, the j -th column in the k -th block is given by

$$q_j^{(k)} = q_{j+(k-1)(n_{fft}-n_{ovlp})+1}. \quad (2.5)$$

Each block is considered as a statistically independent realisation of the flow under the ergodic hypothesis. The motive behind the segmentation of data is to increase the number of ensemble members. In practice, a windowing function is applied to each block to reduce spectral leakage. In this study, we use the symmetric Hamming window

$$w(i+1) = 0.54 - 0.46 \cos\left(\frac{2\pi i}{n_{fft}-1}\right) \quad \text{for } i = 0, 1, \dots, n_{fft}-1. \quad (2.6)$$

Following best practices established by Harris (1978), we only apply windowing for overlapping blocks to avoid excessive loss of information at the boundaries, i.e. if $n_{ovlp} \neq 0$. Subsequently, the weighted temporal discrete Fourier transform,

$$\hat{q}_j^{(k)} = \mathcal{F}\{w(j)q_j^{(k)}\}, \quad (2.7)$$

is performed on each windowed block to obtain the Fourier-transformed data matrix

$$\hat{Q}^{(k)} = [\hat{q}_1^{(k)}, \hat{q}_2^{(k)}, \dots, \hat{q}_{n_{fft}}^{(k)}], \quad (2.8)$$

where $\hat{q}_i^{(k)}$ denotes the k -th Fourier realisation at the i -th discrete frequency. Next, we reorganise the data by frequency. The matrix containing all realisations of the Fourier transform at the l -th frequency reads

$$\hat{Q}_l = [\hat{q}_l^{(1)}, \hat{q}_l^{(2)}, \dots, \hat{q}_l^{(n_{blk})}]. \quad (2.9)$$

From this form, the SPOD modes, Φ , and associated energies, λ , can be computed as the eigenvectors and eigenvalues of the CSD matrix $S_l = \hat{Q}_l \hat{Q}_l^*$. In practice, the number of spatial degrees of freedom, n , is often much larger than number of realisations. In this case, it is more economical to solve the analogous eigenvalue problem

$$\frac{1}{n_{blk}} \hat{Q}_l^* W \hat{Q}_l \Psi_l = \Psi_l \Lambda_l \quad (2.10)$$

for the coefficients Ψ that expand the SPOD modes in terms of the Fourier realisations. In terms of the column matrix $\Psi_l = [\psi_l^{(1)}, \psi_l^{(2)}, \dots, \psi_l^{(n_{blk})}]$, the SPOD modes at the l -th

frequency are recovered as

$$\Phi_l = \frac{1}{\sqrt{n_{blk}}} \hat{Q}_l \Psi_l \Lambda_l^{-1/2}. \tag{2.11}$$

The matrices $\Lambda_l = \text{diag}(\lambda_l^{(1)}, \lambda_l^{(2)}, \dots, \lambda_l^{(n_{blk})})$, where by convention $\lambda_l^{(1)} \geq \lambda_l^{(2)} \geq \dots \geq \lambda_l^{(n_{blk})}$, and $\Phi_l = [\phi_l^{(1)}, \phi_l^{(2)}, \dots, \phi_l^{(n_{blk})}]$ contain the SPOD energies and modes, respectively. By construction, the SPOD modes are orthogonal in the space–time inner product (2.3). At any given frequency, the modes are also orthogonal in the spatial inner product (2.2).

2.2. Data reconstruction

In § 2.2.1, we show how the data can be reconstructed in the frequency domain, that is, the inversion of the SPOD. An alternative approach based on (oblique) projection in the time domain is presented in § 2.2.2. Whether one approach or the other is preferred depends on the specific application. Detailed discussions for each application under consideration in this work can be found in § 3. We will use the term ‘frequency domain’ if the expansion coefficients are computed using the inversion of the SPOD problem, (2.13), and ‘time domain’ if oblique projection, (2.19), is used.

2.2.1. Reconstruction in the frequency domain

The common factor of the different applications of SPOD considered in this study is that they require truncation or re-weighting of the SPOD basis. In practice, this is achieved by modifying the expansion coefficients. The original realisations of the Fourier transform at each frequency can be reconstructed as

$$\hat{Q}_l = \Phi_l \mathbf{A}_l, \tag{2.12}$$

where \mathbf{A}_l is the matrix of expansion coefficients:

$$\mathbf{A}_l = \sqrt{n_{blk}} \Lambda_l^{1/2} \Psi_l^* = \Phi_l^* \mathbf{W} \hat{Q}_l. \tag{2.13}$$

Equation (2.13) shows that the expansion coefficients can either be saved during the computation of SPOD or be recovered later by projecting the Fourier realisations onto the modes. In the following, we omit the frequency index l with the understanding that the SPOD eigenvalue problem is solved at each frequency separately. From (2.12), it can be inferred that each column of the matrix

$$\mathbf{A} = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n_{blk}} \\ a_{21} & a_{22} & \cdots & a_{2n_{blk}} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n_{blk}1} & a_{n_{blk}2} & \cdots & a_{n_{blk}n_{blk}} \end{bmatrix}, \tag{2.14}$$

contains the expansion coefficients that allow for the reconstruction of a specific Fourier realisation from the SPOD modes. *Vice versa*, the coefficients contained in each row of \mathbf{A} can be used to expand a specific SPOD mode in terms of the Fourier realisations. This can most easily be seen by rewriting (2.12) as $\Phi_l = (1/n_{blk}) \hat{Q}_l \mathbf{A}_l^* \Lambda_l^{-1}$.




Figure 1. Schematic representation of overlapping blocks, Hamming window and the SPOD parameters such as n_t , n_{fft} , n_{ovlp} and n_{blk} . The left and right blocks are denoted using blue and red lines, respectively. The vertical line corresponds to the snapshot number 210, and the blue and red circles indicate the corresponding value of the window function.

The Fourier-transformed data of the k -th block can be reconstructed as

$$\hat{\mathbf{Q}}^{(k)} = \left[\left(\sum_i a_{ik} \phi^{(i)} \right)_{l=1}, \left(\sum_i a_{ik} \phi^{(i)} \right)_{l=2}, \dots, \left(\sum_i a_{ik} \phi^{(i)} \right)_{l=n_{fft}} \right]. \quad (2.15)$$

The original data in the k -th blocks $\mathbf{Q}^{(k)}$ can now be recovered using the inverse (weighted) Fourier transform

$$\mathbf{q}_j^{(k)} = \frac{1}{w(j)} \mathcal{F}^{-1} \left\{ \hat{\mathbf{q}}_j^{(k)} \right\}. \quad (2.16)$$

Reconstructing the time series from the reconstructed data segments concludes the inversion of the SPOD.

A schematic of the windowing and blocking strategy is shown in figure 1. The use of overlapping blocks leads to an ambiguity in the reconstruction as the i -th snapshot can either be obtained from the k -th block as $\mathbf{q}_j^{(k)}$, or from the $(k + 1)$ -th block as $\mathbf{q}_{j+n_{ovlp}-n_{fft}}^{(k+1)}$. Different possibilities to remove this ambiguity are described in Appendix A. Based on this discussion, snapshots are reconstructed as averages of two reconstructions from overlapping blocks, weighted by the relative value of their windowing function. Partial reconstructions in the frequency domain are readily achieved by zeroing the expansion coefficients of specific modes prior to applying the inverse Fourier transform.

2.2.2. Reconstruction in the time domain

As an alternative to the reconstruction in the frequency domain, we present in the following a projection-based approach in the time domain. This approach is computationally efficient, and has the advantage that it can be applied to new data that was not used to compute SPOD and to individual snapshots. However, the time-domain reconstruction does not leverage the orthogonality of the SPOD modes in the space–time inner product. Instead, it is based on an oblique projection of the data onto the modal basis. We start by representing the data as a linear combination of the SPOD modes as

$$\mathbf{Q} \approx \tilde{\Phi} \tilde{\mathbf{A}}. \quad (2.17)$$

The matrix $\tilde{\Phi}$ contains the basis of SPOD modes at all frequencies. Arranging the basis vectors by frequency first, we write

$$\tilde{\Phi} = \left[\phi_1^{(1)}, \phi_1^{(2)}, \dots, \phi_1^{(n_{blk})}, \phi_2^{(1)}, \phi_2^{(2)}, \dots, \phi_2^{(n_{blk})}, \dots, \phi_{n_{fft}}^{(1)}, \phi_{n_{fft}}^{(2)}, \dots, \phi_{n_{fft}}^{(n_{blk})} \right]. \quad (2.18)$$

Assuming that $\tilde{\Phi}$ has full column rank, the matrix of expansion coefficient is obtained from the weighted oblique projection

$$\tilde{\mathbf{A}} = \left(\tilde{\Phi}^* \mathbf{W} \tilde{\Phi} \right)^{-1} \tilde{\Phi}^* \mathbf{W} \mathbf{Q}. \tag{2.19}$$

The oblique projection is required as SPOD modes at different frequencies are not orthogonal in the purely spatial inner product, $\langle \cdot, \cdot \rangle_x$, defined in (2.2). We furthermore use a weighted oblique projection based on the weight matrix \mathbf{W} to guarantee compatibility with this inner product. Using the oblique projection, a single snapshot $\mathbf{q} = \mathbf{q}(\mathbf{x}, t)$ is represented in the SPOD basis as

$$\tilde{\mathbf{q}} = \tilde{\Phi} \left(\tilde{\Phi}^* \mathbf{W} \tilde{\Phi} \right)^{-1} \tilde{\Phi}^* \mathbf{W} \mathbf{q}, \tag{2.20}$$

and the entire data are recovered as

$$\tilde{\mathbf{Q}} = \tilde{\Phi} \tilde{\mathbf{A}}. \tag{2.21}$$

In the case where $\tilde{\Phi}$ is rank-deficient or ill-conditioned, $\tilde{\Phi}^* \mathbf{W} \tilde{\Phi}$ is not invertible. For a general rank- r matrix, we perform the symmetric eigenvalue decomposition

$$\tilde{\Phi}^* \mathbf{W} \tilde{\Phi} = \mathbf{U} \mathbf{D} \mathbf{U}^*, \quad \text{with } \mathbf{U} = [\mathbf{U}_1 \ \mathbf{U}_2], \ \mathbf{D} = \begin{bmatrix} \mathbf{D}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}, \tag{2.22}$$

where $\mathbf{D}_1 = \text{diag}(d_1, d_2, \dots, d_r)$ with $d_1 \geq d_2 \geq \dots \geq d_r > 0$ is the diagonal matrix of (numerically) non-zero eigenvalues and \mathbf{U}_1 is the corresponding matrix of orthonormal eigenvectors. In some applications, we desire a more aggressive truncation to rank $k < r$. For full-rank reconstructions, we use a truncation threshold of $d_k/d_1 = 10^{-6}$ in this work. Denoted by

$$\tilde{\mathbf{A}}_{\{k\}} = \mathbf{U}_{\{k\}} \mathbf{D}_{\{k\}}^{-1} \mathbf{U}_{\{k\}}^* \tilde{\Phi}^* \mathbf{W} \mathbf{Q}, \tag{2.23}$$

is the rank- k approximation of $\tilde{\mathbf{A}}$, where $\mathbf{D}_{\{k\}} = \text{diag}(d_1, d_2, \dots, d_k)$ and $\mathbf{U}_{\{k\}} = [u_1, u_2, \dots, u_k]$. In the truncated basis, $\mathbf{U}_{\{k\}} \mathbf{D}_{\{k\}}^{-1} \mathbf{U}_{\{k\}}^*$ approximates $\left(\tilde{\Phi}^* \mathbf{W} \tilde{\Phi} \right)^{-1}$. In the following, we demonstrate how the above and other truncation and partial reconstruction strategies can be used to achieve a number of objectives in the processing of flow data.

3. Applications of SPOD: low-rank reconstruction, denoising, frequency–time analysis and prewhitening

In this section, four different uses of SPOD-based applications are introduced and demonstrated for the example of a turbulent jet. The theoretical background of each application is presented in the context of the jet. In particular, low-dimensional reconstruction is discussed in § 3.1, denoising in § 3.2, frequency–time analysis in § 3.3 and prewhitening in § 3.4.

We consider the large-eddy simulation (LES) data of an isothermal subsonic turbulent jet, the Reynolds number, Mach number and temperature ratio are defined as $Re = \rho_j U_j D / \mu_j = 0.45 \times 10^6$, $M_j = U_j / c_j = 0.4$ and $T_j / T_\infty = 1.0$, respectively, where ρ is the density, U velocity, D nozzle diameter, μ dynamic viscosity, c speed of sound and T temperature. The subscripts j and ∞ refer to the jet inlet and free-stream conditions, respectively. We use the LES data computed by Brès & Lele (2019). The simulation

was performed using the compressible flow solver ‘Charles’ (Brès *et al.* 2017) on an unstructured grid using a finite-volume method. The reader is referred to Brès *et al.* (2018); Brès & Lele (2019) for further details on the numerical method. The LES database consists of 10 000 snapshots sampled at an interval of $\Delta t c_\infty / D = 0.2$ acoustic time units. Data interpolated on a cylindrical grid spanning $x, r \in [0, 30] \times [0, 6]$ was used in this analysis. The flow is non-dimensionalised by the nozzle exit values, namely velocity by U_j , pressure by $\rho_j U_j^2$, length by the nozzle diameter D and time by D/U_j . Frequencies are reported in terms of the Strouhal number $St = fD/U_j$. For simplicity, and without loss of generality, we perform our analysis only on the pressure field in what follows. Refer to Freund & Colonius (2009) for analysis on different energy norms and POD modes of a jet. We further exploit the rotational symmetry of the jet and consider individual azimuthal Fourier components. The helical ($m = 1$) component provides an example of complex data.

To determine the spectral estimation parameters for the SPOD, we follow the guidelines provided in Schmidt & Colonius (2020) and Schmidt *et al.* (2018), which, in turn, follow standard practice in spectral estimation. The SPOD is computed for blocks containing $n_{fft} = 256$ snapshots with 50 % overlap, resulting in a total number of $n_{blk} = 77$ blocks. A 50 % overlap is used to minimise the variance of the spectral estimate (Welch 1967).

Since SPOD yields one set of eigenpairs per frequency, we may investigate the contributions of different frequencies independently. The SPOD eigenvalues are represented in the form of a spectrum, reminiscent of a power spectrum, in figure 2(a). Grey lines of decreasing intensity connect eigenvalues of constant mode number and decreasing mode energy. The first, most energetic mode is shown in magenta. The red line represents the sum of all eigenvalues and corresponds to the power spectral density (PSD) integrated over the physical domain. This line of the integral total energy can be compared with truncated sums of eigenvalues, that is, the energy contained in reconstructions of different ranks. As the eigenvalues are sorted by energy, the lines corresponding to the truncated sums of the leading 3 and 10 eigenvalues fall between the leading-eigenvalue spectrum and the total energy curve. Figure 2(b) shows the normalised cumulative energy content, independent of frequency. The first and the leading 10 modes contain 30 % and 80 % of the total energy, respectively. Figure 2(c) shows the percentage of energy accounted for by each mode as a function of frequency. At low frequencies, the first few modes contain a high percentage of energy, whereas the energy is more dispersed at higher frequencies. The solid and dashed white lines indicate the number of modes required to retain 50 % and 90 % of the total energy at each frequency.

The first and second modes at two representative frequencies are shown in figure 3. The frequency $St = 0.5$ corresponds to the maximum difference between first and second eigenvalues. The leading mode at this frequency shows a Kelvin–Helmholtz (KH) wavepacket (Suzuki & Colonius 2006; Gudmundsson & Colonius 2011; Schmidt *et al.* 2018) in the shear layer of the jet. A similar, but a more compact wavepacket structure is observed at $St = 1.0$. The suboptimal mode at both frequencies exhibit a multi-lobed wavepacket structure, whose amplitude peaks near the end of the potential core at $x \approx 6$. The reader is referred to Schmidt *et al.* (2018) and Tissot *et al.* (2017) for a physical discussion of this observation and the link to non-modal instability. In the present context, our preliminary interest is in the desirable mathematical property of SPOD that guarantees that the modes optimally represent the turbulent flow field in terms of the space–time inner product, (2.3). In the following, different uses of low-dimensional reconstructions that use SPOD modes as basis vectors are introduced and discussed.




Figure 2. The SPOD spectra of the turbulent jet for $m = 1$. All eigenvalues (grey lines) and the sum of all eigenvalues (red line), corresponding to the integral PSD, are shown. The normalised cumulative energy content and the percentage of energy accounted by each mode as a function of frequency are shown in (b,c), respectively. The solid and dashed white lines indicate the number of modes required to retain 90 % and 50 % energy at each frequency.




Figure 3. The SPOD modes at $St = 0.5$ (a,b) and at $St = 1.0$ (c,d). The leading modes are shown in (a,c) and the suboptimal modes are shown in (b,d). (a) SPOD mode 1, $St = 0.50$, (b) SPOD mode 2, $St = 0.50$, (c) SPOD mode 1, $St = 1.00$ and (d) SPOD mode 2, $St = 1.00$.

3.1. Low-dimensional flow-field reconstruction

Since SPOD seeks an optimal series expansion for each frequency, the choice of what eigenpairs to include in a low-dimensional reconstruction is not obvious. Here we first discuss the most elementary way of truncation based on the frequency-wise optimality property, that is, a certain number of modes is retained at each frequency. We refer to this as a $n_{modes} \times n_{freq}$ -mode reconstruction, where n_{modes} is the number of modes retained at each frequency, and $n_{freq} = n_{fft}/2 + 1$ is the number of positive frequencies,

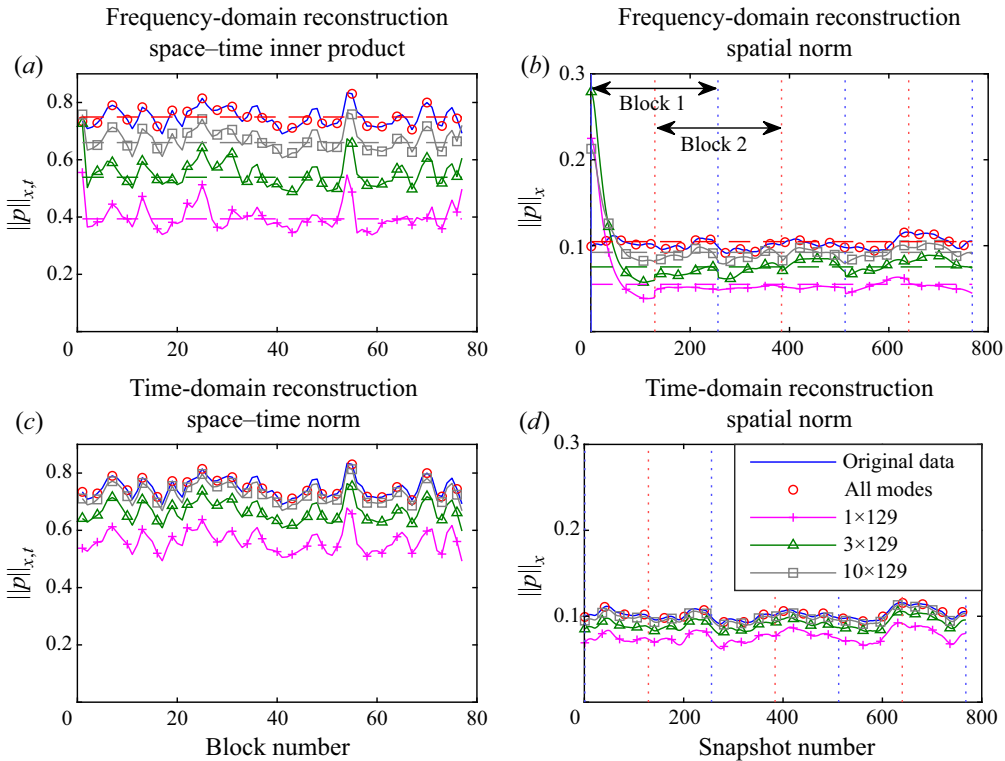


Figure 4. Low-dimensional reconstruction of the jet data: frequency-domain reconstruction (a,b), and time-domain reconstruction (c,d) in terms of the space-time norm (a,c) and the spatial norm (b,d). The original data (blue lines), shown for reference, are compared with the full reconstructions using all modes, and reconstructions using 10×129 , 3×129 and 1×129 modes. Summed SPOD mode energies are shown as dashed lines. Vertical dotted blue and red lines in (b,d) indicate the left and right blocks in the reconstruction (see figure 1).

including zero. If all modes are linearly independent, then the overall rank of the reconstructions is given by the total number of basis vectors, $n_{modes}n_{freq}$.

Following the discussion in § 2, we present two means of obtaining an SPOD-based low-dimensional reconstruction:

- (i) in the frequency domain (see § 2.2.1) using (2.16); and
- (ii) in the time domain (see § 2.2.2) using (2.20).

The frequency-domain approach directly follows from the mathematical definition of SPOD, and was previously used by Citriniti & George (2000), Jung *et al.* (2004), Johansson & George (2006b) and Tinney *et al.* (2008a). The time-domain approach can be viewed as the most general approach that can be applied to any given modal basis. It is not specific to SPOD, but commonly used for low-order modelling.

In the following, we first compare different low-dimensional reconstructions in terms of their block-wise and snapshot-wise energy in figure 4. It follows from (2.3), that the energy of a single block is

$$\|q\|_{x,t}^2 = \langle q, q \rangle_{x,t} = \int_{\Delta T} \int_{\Omega} q^*(x', t) \mathbf{W}(x') q(x', t) dx' dt, \quad (3.1)$$




Figure 5. Low dimensional reconstruction in the frequency-domain with truncation correction: (a) space–time norm; (b) spatial norm. Panels (a,b) correspond to figure 4(a,b), respectively, but with a correction for the truncated modes. The correction is facilitated by adding the energies of the truncated modes (given by their SPOD eigenvalues).

where $\Delta T = [t_{1+(k-1)(n_{ffl}-n_{ovlp})}, t_{n_{ffl}+(k-1)(n_{ffl}-n_{ovlp})}]$ is the time interval of the k -th block. The spatial norm (2.2) measures the energy present in each snapshot. Both the block-wise and snapshot-wise pressure norms are computed for both reconstruction approaches. The evolution of the space–time norm is shown for the entire database in figure 4(a,c). For the spatial norm, we focus on the first 768 snapshots in figure 4(b,d). This segment corresponds to five blocks and exhibits dynamics that is representative of the rest of the data.

Low-dimensional reconstructions using 1×129 , 3×129 , 10×129 modes and the reconstruction using all modes are presented in figure 4. The dimension of the modal bases directly reflects their ability to capture the pressure norms of the data. The full-dimensional reconstructions in the frequency and time domain recover the data completely. Notably, the dynamics in space–time norm is accurately captured, even by the 1×129 -mode reconstruction. For a fixed number of modes, the time domain approach captures more energy and provides a better approximation of the data than the frequency-domain approach. Take as an example the 10×129 basis: the time-domain reconstruction accurately approximates for the full data (figure 4c,d), which is notably underpredicted by frequency-domain reconstruction (figure 4a,b). This difference can be understood by considering the SPOD energy content of the reconstruction. The dashed lines in figure 4(a,b) denote this energy, which is given by the sum of the first n_{modes} eigenvalues over all frequencies. As expected, the space–time and spatial norm of the different reconstructions fluctuate about the sum of the eigenvalues. Higher energies are obtained by the time-domain reconstruction in figure 4(c,d) as the modal expansion coefficients obtained via oblique projection are not bound to specific frequencies. In what follows, we will see again and again that this flexibility of the expansion coefficients of the time-domain approach leads to an overall better reconstruction. This additional degree of freedom can be leveraged to obtain an accurate reconstruction of the flow dynamics. It is, in fact, the optimality property of the oblique projection, (2.20), that guarantees that the time-domain approach yields the best possible approximation in a least-square sense. For example, the 1×129 time-domain reconstruction seen in figure 4(d) is sufficient to capture the dynamics of the original data accurately. From figure 4(b), it can be seen that the frequency-domain reconstruction significantly overpredicts the pressure norm during the first 32 snapshots. This effect only occurs for the two outermost blocks, which do

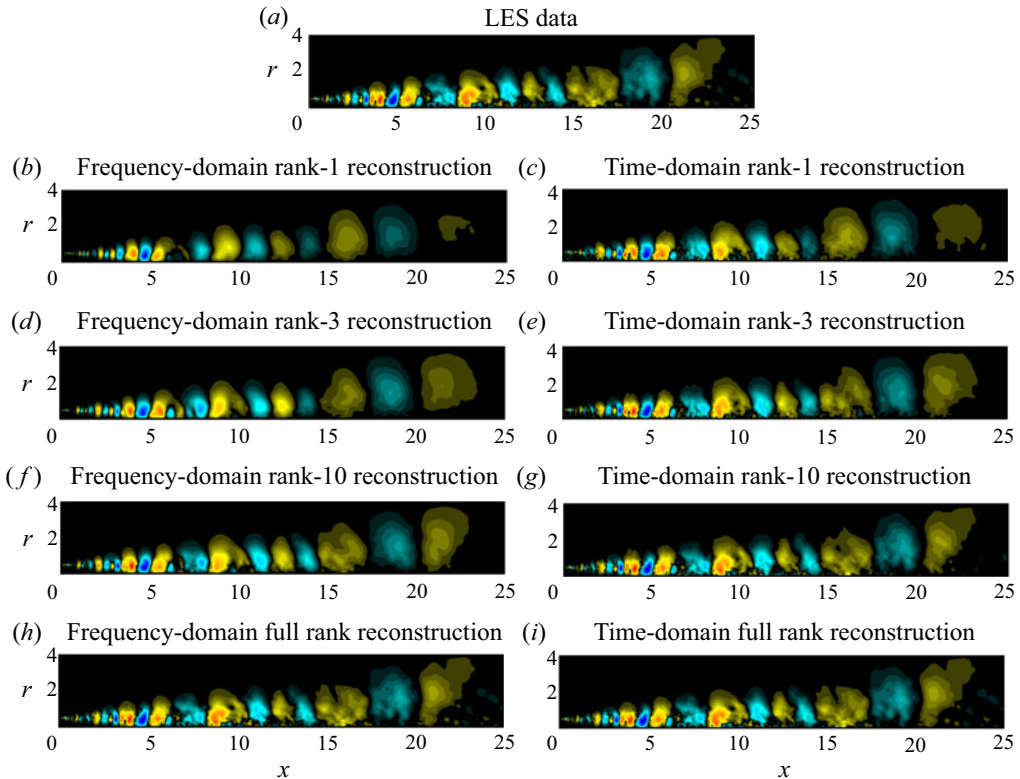


Figure 6. Instantaneous pressure field: (a) original flow field is shown; (b,d,f,h) reconstructions in the frequency domain; (c,e,g,i) reconstructions in the time domain. Flow fields reconstructed using 1×129 modes, 3×129 modes, 10×129 modes and all modes are shown in (b,c), (d,e), (f,g) and (h,i), respectively. Contours in (a–i) are reported on the same colour axis.

not possess neighbouring blocks in one direction. In [Appendix A \(figure 20\)](#), we show by comparison with a rectangular window that this error is a result of the Hamming window. The presence of the windowing effect in the first and last blocks is equally reflected in the space–time norm, see [figure 4\(a\)](#).

[Figure 5](#) demonstrates that the frequency-domain approach accurately recovers the mode energies given by the SPOD eigenvalues. By adding the residual energy contained in the truncated eigenvalues, both the space–time norm ([figure 5a](#)) and the spatial norm ([figure 5b](#)) can be collapsed to the total energy. The remaining differences are largely due to the windowing effect.

[Figure 6](#) compares a single time instant of the original data in (a) with reconstructions of increasing fidelity in the frequency (left) and the time domain (right). Both the 1×129 -mode reconstructions shown in [figure 6\(b,c\)](#) capture the dominant wavepackets. However, the frequency-domain reconstruction lacks the detail of the time-domain reconstruction. The higher accuracy of the time-domain reconstruction can be explained by its less stringent nature. As the leading SPOD modes often represent a spatially highly confined structure, other structures associated with the same frequency cannot be represented by the frequency-domain reconstruction. The KH-type wavepacket seen in [figure 3\(c\)](#) is a good example of such a confined structure. This difference between the approaches also explains the better reconstruction of the integral energy in the

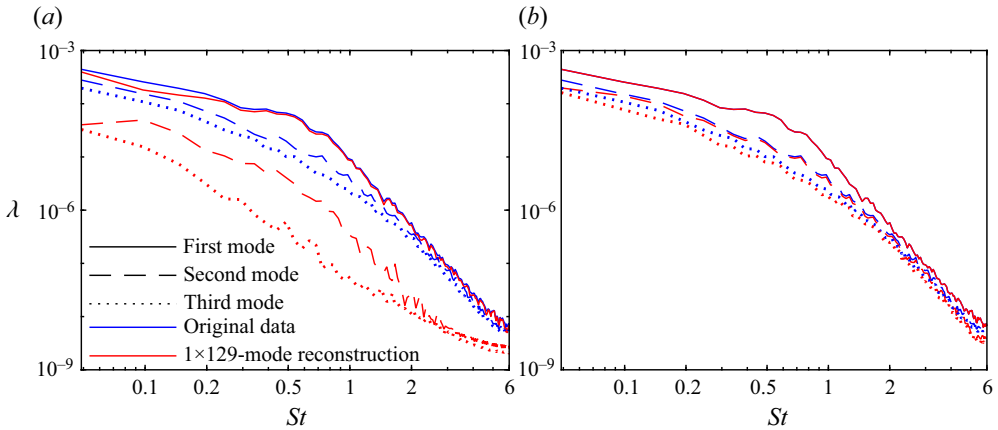


Figure 7. Comparison between SPOD eigenvalue spectra of the original data (blue lines) and 1×129 -mode reconstructions (red lines): (a) frequency domain; (b) time domain. Solid, dashed and dotted lines denote the first, second and third modes, respectively.




Figure 8. The 2-norm error of the pressure field reconstructed with different rank approximation: (red lines with squares) reconstruction in the frequency domain; (blue lines with circles) reconstruction in the time domain.

time domain, as previously observed in figure 4. The higher-dimensional versions for both approaches shown in figure 6(d–g) yield increasingly more detailed and accurate reconstructions. Both approaches yield reconstructions that are indistinguishable from the original data when all modes are used (see figure 6h,i).

We infer from figures 4–6 that reconstruction in the time domain provides a better estimate of the flow field than the frequency-domain version. To understand this observation, figure 7 reports the SPOD eigenspectra of the 1×129 -mode reconstructions and compares them to those of the full data. Only the leading three eigenvalues are shown for clarity. The leading eigenvalue of the frequency-domain reconstruction in figure 7(a) approximately follows the full data with some discrepancies at lower frequencies. No such discrepancies are observed for the time-domain reconstruction in figure 7(b); in fact, the leading eigenvalue spectra are indistinguishable. Contrast this observation with the expectation that a $1 \times n_{freq}$ frequency-domain reconstruction should exactly reproduce

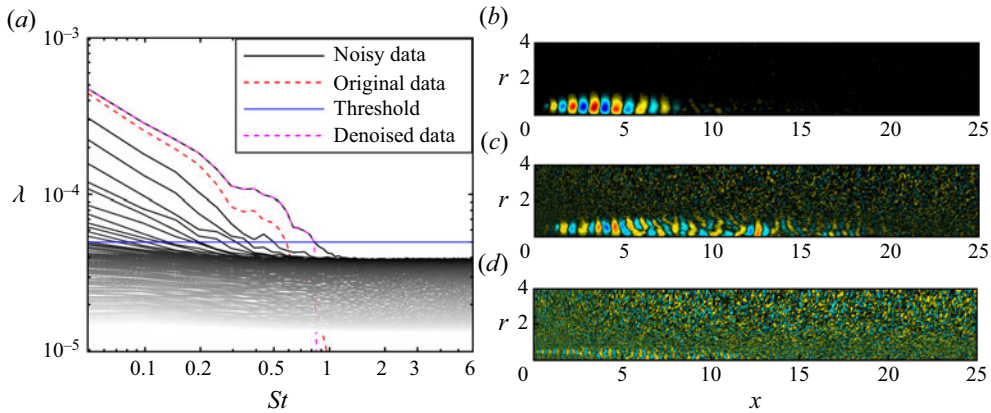


Figure 9. The SPOD of data subjected to additive Gaussian white noise: (a) SPOD spectrum (black lines), leading SPOD eigenvalue of the original data (red dashed line), threshold of 5×10^{-5} (blue line) and the leading SPOD eigenvalue of the denoised data (magenta dashed line); (b) leading SPOD mode at $St = 0.5$; (c) fifth SPOD mode at $St = 0.5$; (d) leading SPOD mode at $St = 1.5$.

the leading-mode eigenspectrum, and that all higher-eigenvalue spectra are expected to be zero. In the context of figure 21 in Appendix A, we show that this is an effect of windowing that is not observed when using a rectangular window. The time-domain reconstruction, on the other hand, accurately approximates the first, and, to some degree, the leading suboptimal eigenvalue spectra. This again demonstrates the higher accuracy of the time-domain reconstruction that results from the higher flexibility of the expansion.

To quantify the accuracy of the two approaches, figure 8 compares their 2-norm errors as a function of the number of basis vectors (modes). For both the methods, the error reduces significantly as the number of modes retained per frequency increases from one to ten. For a fixed number of modes, the time-domain reconstruction is consistently more accurate. Recall that this is guaranteed by the optimality property of the oblique projection. It is, in fact, observed that the error of the time-domain reconstruction approaches machine precision for 40 or more modes per frequency. For the frequency-domain approach this only occurs for the full reconstruction using all modes.

3.2. Denoising

After using SPOD truncation for the low-dimensional approximation previously described, we now explore its potential for denoising. We will show that additive noise is captured by certain parts of the spectrum, and that truncation of these parts leads to efficient denoising. This strategy is most efficiently implemented in the frequency domain. The local-in-time optimality of the time-domain approach is a hindrance in this context as it tends to reconstruct the noise. However, the one-to-one correspondence between modes and frequencies of the frequency-domain approach leads to efficient denoising.

We demonstrate denoising on additive Gaussian white noise, arguably the most common type of noise occurring in experimental environments. In particular, we add Gaussian white noise that has a standard deviation equal to the spatial mean of the standard deviation along the lipline of the pressure data. This scenario is very similar, for example, to heavily contaminated PIV data in which the variance of the noise is of the same order as the variance of the physical phenomena of interest. The SPOD eigenvalue spectrum of this noisy data is shown in figure 9(a). Most noticeably, the addition of noise has introduced a

noise floor at $\lambda \approx 4 \times 10^{-5}$, effectively cutting off the spectrum at $St \approx 1.5$. Information above this frequency lies below the noise floor and is not directly accessible. The leading eigenvalue of the original data (red dashed line) is shown for comparison. It lies well below the leading eigenvalue of the noisy data, which is elevated owing to the energy contained in the added noise. To illustrate the ability of SPOD to differentiate between spatially correlated, physical structures and noise, examples of modes that are above and below the noise floor are compared in figure 9(b–d). The leading mode at $St = 0.5$ (figure 9b) clearly reveals the KH wavepacket and is indistinguishable from the corresponding mode of the original data shown in figure 3(a). The fifth mode at $St = 0.5$ (figure 9c) and the leading mode at $St = 1.5$ (figure 9d), by contrast, are heavily contaminated by noise. The noise floor in the SPOD spectrum is found to be a very good indicator of this distinction. We therefore propose a denoising strategy based on hard thresholding of the spectrum. In this example, we pick a threshold of 5×10^{-5} (blue line), slightly above the noise floor. To address the effect of the truncation on the SPOD spectrum, we report the leading SPOD eigenvalue of the denoised data (magenta dashed line) in the same figure. It coincides with the leading eigenvalue of the noisy data up to the point where it intersects with the threshold limit, beyond which it falls off sharply, giving it the characteristics of a low-pass filter. A closer analysis of the truncated and original spectra reveals that the denoised field contains only 2.6% of the energy of the noisy field, but that it contains 92.7% of the energy of the original flow field. Another positive side effect is that only 74 out of 9933 modes ($n_{freq} \times n_{blk}$) have been retained, resulting in a space saving of 99.26%. These significant space savings are an advantage over standard denoising strategies based on low-pass filtering.

Representative instantaneous snapshots of the original data are compared with their noisy and denoised counterparts, and to the result of standard low-pass filtering in figure 10. The standard low-pass filter uses the cut-off frequency of $St = 0.8$ of the SPOD approach (inferred from figure 9) in the truncation of the long-time Fourier transform. A higher threshold, and its associated cut-off frequency, was found to lead to more aggressive filtering that can partially remove relevant flow structures. A threshold below the noise floor, on the other hand, leads to unsatisfactory noise rejection. In practice, a good trade-off between noise rejection and preservation of physically relevant flow structures is achieved by using the SPOD spectrum as a gauge to choose a threshold slightly above the noise floor. A comparison of the denoised data in figure 10(c,d) with the noisy data in figure 10(b) shows that significant noise reduction was achieved in all parts of the domain using both strategies. The resulting denoised flow fields clearly reveal the flow structures present in the original data. By visual inspection of the filtered pressure fields shown in figures 10(c) and 10(d), the SPOD-based strategy appears somewhat more efficient at removing the noise.

For a more quantitative assessment, we compare the denoised flow fields in terms of two quantities. First, their signal-to-noise ratio (SNR) along the lipline, and second, the relative error between the denoised snapshots and the original data. The SNR is defined as

$$\text{SNR} = \frac{P_{\text{signal}}}{P_{\text{noise}}} = \frac{\sigma_{\text{signal}}^2}{\sigma_{\text{noise}}^2}, \tag{3.2}$$

where P is power and σ standard deviation. We further define the integral (over the physical domain) error as

$$\text{error} = \frac{\|q - \check{q}\|_x}{\|q\|_x}, \tag{3.3}$$




Figure 10. Comparison of noisy and denoised instantaneous pressure fields: (a) original LES data; (b) LES data with additive Gaussian white noise; (c) SPOD-based denoised flow field; (d) low-pass filtered flow field. Denoising is achieved by rejecting all SPOD eigenpairs with $\lambda < 5 \times 10^{-5}$. This hard threshold is indicated in figure 9.




Figure 11. Comparison of the two denoising strategies: (a) the SNR along the lipline ($r = 0.5$) for the noise added flow field (black line), SPOD-based denoised flow field (red lines), and the low-pass filtered flow field (blue line); (b) error of the noisy data, SPOD-based denoised data, and the low-pass filtered data. The amplitude of the additive noise was adjusted such that the average SNR along the lipline is one.

where q and \check{q} are the original and the denoised flow fields, respectively. Figure 11(a) compares the SNR along the lipline at $r = 0.5$ for the noisy and the denoised flow fields. Both methods achieve an increase in the SNR over large parts of the domain. The low-pass filter performs better for $x \lesssim 10$, and the SPOD-based filter beyond that point. For $x \lesssim 2.5$, the SPOD-filtered pressure field exhibits a marginally lower SNR than the unfiltered data. We find that this is the result of the aggressive truncation of the high-frequency components by the SPOD-based filter. A result that is almost identical to that of the low-pass filter can be achieved by lowering the λ -threshold (a similar value for both methods is used here for consistency). Figure 11(b) compares the time traces of the errors of the noisy and the two denoised flow fields. The error of the noisy data serves as a reference, and it is observed that both methods significantly reduce this error. The SPOD-based approach performs consistently better than the low-pass filter. This result is consistent with the visual observation of the denoised fields in figure 10(c,d). Note that the threshold is an adjustable parameter in both methods. In practice, we find that by adjusting this parameter qualitatively very similar results can be obtained by both methods. This, however, leaves the SPOD-based approach with the advantage of significant data reduction.




Figure 12. The SPOD-based frequency–time diagrams obtained using the time-domain approach: (a) first ten modes at each frequency; (b) leading mode at each frequency. The SPOD of the pressure field is considered.

3.3. Frequency–time analysis

Intermittency, that is, the occurrence of flow events at irregular intervals, is an inherent feature of any turbulent flow. A common approach for the characterisation of intermittent behaviour is frequency–time analysis. Arguably the most commonly used tools of frequency–time analysis are WT and STFT. Their outcomes are scalograms and spectrograms, respectively, which indicate the presence of certain scales (WT), or frequency components (STFT), at certain times. Both methods are signal-processing techniques that are applied to one-dimensional time series, and therefore only quantify intermittency locally. As an alternative to this local perspective, we demonstrate how SPOD expansion coefficients can be used to study the intermittency of the spatially coherent flow structures represented by the modes. Next, frequency–time analyses based on both time-domain and frequency-domain reconstructions are introduced and compared.

3.3.1. Time-domain approach

We first consider the time-domain approach, in which the expansion coefficients obtained via oblique projection readily describe the temporal behaviour of each mode. The amplitudes of the expansion coefficients computed from (2.19), $|\sum_{j=1}^{n_{modes}} \tilde{a}^{(j)}(f_t, t)|$, hence, yield the desired frequency–time representation for the leading n_{modes} . The expansion coefficients are calculated using the full basis, i.e. $\tilde{\phi}$ in (2.18), consists of all modes at all frequencies. Subsequently, we only consider the expansion coefficients of the leading n_{modes} modes at each frequency. An alternative approach is to perform the oblique projection using a reduced basis that consists of only the leading n_{modes} modes at each frequency. We find that the first approach is preferable in the context of frequency–time analysis and is explained in the Appendix B (see figure 22). The frequency–time diagrams for $n_{modes} = 10$, and 1, are shown in figures 12(a) and 12(b), respectively. The leading 10 modes correspond approximately to 80 % of the total energy as shown in figure 2(b). Most of the energy is concentrated at low frequencies, $St \lesssim 0.2$, as expected from the eigenvalue spectrum in figure 2. The eigenvalue spectrum provides a statistical representation of the structures that are coherent in space and time, whereas the frequency–time diagrams provide a temporal information of these structures. Bright yellow spots indicate high similarity of the instantaneous flow field with the leading mode in (b). These regions also correspond to high-energy events as we will show in figure 14.




Figure 13. The SPOD-based frequency–time diagrams obtained using the convolution approach: (a) first ten modes at each frequency; (b) leading mode at each frequency. The SPOD of the pressure field is considered. The cyan box in (b) is centred around the global maximum of the instantaneous energy at $t = 1380$, analysed in figures 14 and 15(c,d) below.

3.3.2. Convolution-based approach

A direct way of using SPOD for frequency–time analyses is in terms of the SPOD expansion coefficients. Since each block is associated with a finite time interval, this approach requires the computation of SPOD using an overlap of $n_{ovlp} = n_{fft} - 1$ to obtain time-resolved coefficients (Towne & Liu 2019). This approach assumes that the value of the expansion coefficient obtained from a finite time segment (block) represents the instantaneous frequency content at the centre of the time segment. A limitation of this approach is its high memory requirement (3.6 TB for the present example). As an alternative, we propose a computationally tractable way of calculating time-continuous expansion coefficients based on the convolution theorem. Applying the convolution theorem to the inverse SPOD problem yields

$$\begin{aligned} a_l^{(i)}(t) &= \left(\phi_l^{(i)} \exp(-i2\pi f_l t) \otimes q \right) (t) \\ &= \int_{\Delta T} \int_{\Omega} \left(\phi_l^{(i)}(x) \right)^* W(x) q(x, t + \tau) w(\tau) \exp(-i2\pi f_l \tau) dx d\tau, \end{aligned} \quad (3.4)$$

where \otimes indicates the convolution between the time evolving SPOD mode and the data, which takes into the account the windowing function, $w(\tau)$, and the weight matrix, W . In practice this convolution is computed by expanding the SPOD mode in time as $\phi_l^{(i)} \exp(-i2\pi f_l t)$ and convolving it over the data one snapshot at a time. In this step, we leverage the orthogonality property of the SPOD mode in the space–time inner product, which allows us to compute the expansion coefficient one at a time. If the SPOD was computed using an overlap of $n_{fft}-1$, then the expansion coefficients, $a_l^{(i)}$, obtained from (3.4) and (2.13) are mathematically identical. Here, the underlying idea is to apply the continuously discrete convolution integral to the SPOD mode computed with a significantly lower overlap to make it computationally feasible. We confirmed that the frequency–time diagrams of the expansion coefficients obtained using (3.4) for an overlap of 50% are virtually indistinguishable to those obtained from (2.13) for an $n_{ovlp} = n_{fft} - 1$ (shown in the Appendix C, figure 24). This is to be expected since the convergence of the SPOD modes does not improve significantly for overlap over 50%. In practice, the convolution integral in (3.4) is most efficiently computed using fast-Fourier transforms.

Frequency–time diagrams of the expansion coefficients for the convolution approach are shown in figure 13. Figure 13(a) and 13(b) show the contribution of the leading




Figure 14. Temporal evolution of the pressure 2-norm in the vicinity of its global maximum at $t = 1380$. The convolution approach, based on (3.4), is compared to the time- and frequency-domain approaches, previously shown in figure 4(b,d). The 1×129 -mode reconstructions are shown.

ten modes and the leading mode, respectively. These frequency–time diagrams appear less detailed compared to the frequency–time diagrams of the time-domain approach. In the context of figure 15, though, we will show that figures 12 and 13 basically contain the same information. For now, it is sufficient to note that the convolution and time-domain approaches detect the same trends. Take as an example, the high energy events occurring at low frequency in the time ranges, $550 \lesssim t \lesssim 590$ and $1370 \lesssim t \lesssim 1400$, in the spectrograms of figures 12(b) and 13(b).

3.3.3. Comparison of methods and interpretation of results

Next, we investigate if the high similarity between the instantaneous flow field and the modes indicates high energy. Figure 14 shows the temporal evolution of the pressure 2-norm for the original data, low-dimensional 1×129 -mode reconstructions using the time-domain, frequency-domain and the convolution approaches. The pressure 2-norm of the data is highly underpredicted by the time-domain approach that uses a full basis ($n_{blk} \times 129$), but is able to capture the major trends of the original data. The time-domain reconstruction performed using a modal basis of 1×129 is also shown for comparison. This accurately follows the spatial norm of the original data, except for an offset, similar to figure 4(d). The frequency-domain curve also shows a similar trend. In addition, the spatial norm of the 1×129 -mode reconstruction using the convolution approach is shown. It follows the trend of the original data and attains its global maximum at the same time instant. Minor differences between the convolution (time-continuous) and frequency-domain (50 % overlap) approaches are expected, see Appendix C. All curves in figure 14 peak at the time of maximum instantaneous energy, previously indicated in figure 13(b). This indicates that a high similarity between the instantaneous flow field and the leading mode implies high overall energy. We highlight that this finding is not self-evident as the leading SPOD mode represents the most energetic flow structure in a purely statistical sense. The important physical insight is that the intermittent occurrence of large-scale coherent structures is directly associated with high-energy events.

To understand the qualitative differences of the frequency–time diagrams in figures 12 and 13, we now look at the expansion coefficients of the two approaches. Figure 15

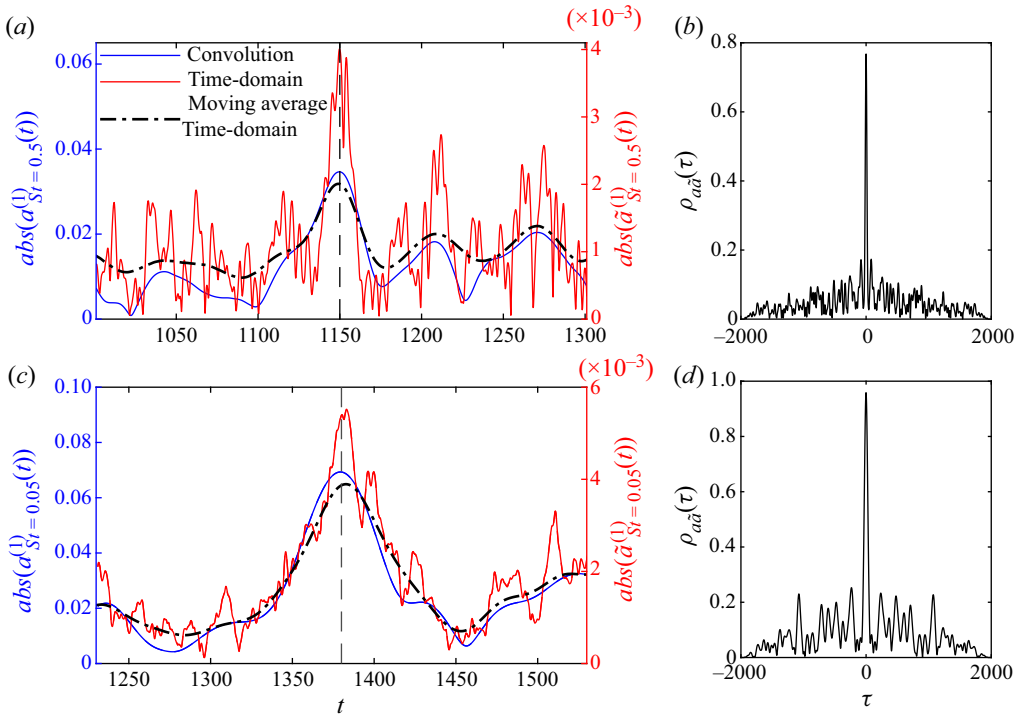


Figure 15. Oblique projection and convolution-based expansion coefficients of the leading mode at $St = 0.5$ (top) and $St = 0.05$ (bottom): (a,c) time traces in the vicinity of their global maximum (black dashed line); (b,d) cross-correlation of the expansion coefficients obtained using the two different approaches. The black dash-dotted line denotes the weighted moving mean of the time-domain expansion coefficient. The moving mean uses as weights the same Hamming windowing function as the SPOD.

compares the expansion coefficient of the leading mode, obtained by the time-domain and convolution approaches. As an example, the expansion coefficient at $St \approx 0.50$ is shown in figure 15(a). It is centred around its global maximum, in the time interval $1000 \leq t \leq 1300$. We observe that the time traces of the expansion coefficients obtained from the two approaches show similar trends. In particular, the local peaks occur at similar locations, with both curves exhibiting the global maximum at $t = 1150$ (black dashed line). Compared to the time-domain approach, the convolution curve is smoother and resembles a moving average of the time-domain curve (black dash-dotted line). For optimal comparison with the convolution approach, the moving average is computed by weighting the time-domain curve by the Hamming window in (2.6) and averaging over 256 points. To further quantify the relation between the expansion coefficients of the two approaches, we show the cross-correlation coefficient in figure 15(b). The cross-correlation coefficient confirms the observation that the expansion coefficients obtained from the two approaches are similar, by demonstrating a cross-correlation coefficient of 0.77 at 0-time lag ($\tau = 0$). We have confirmed that this correspondence holds in general. In figures 15(c) and 15(d), for example, the same trends are observed for the expansion coefficients at the lower frequency of $St = 0.05$ over the time interval previously shown in figure 14. From figure 15, we infer that the intermittency of the coherent structures can be captured using both approaches.

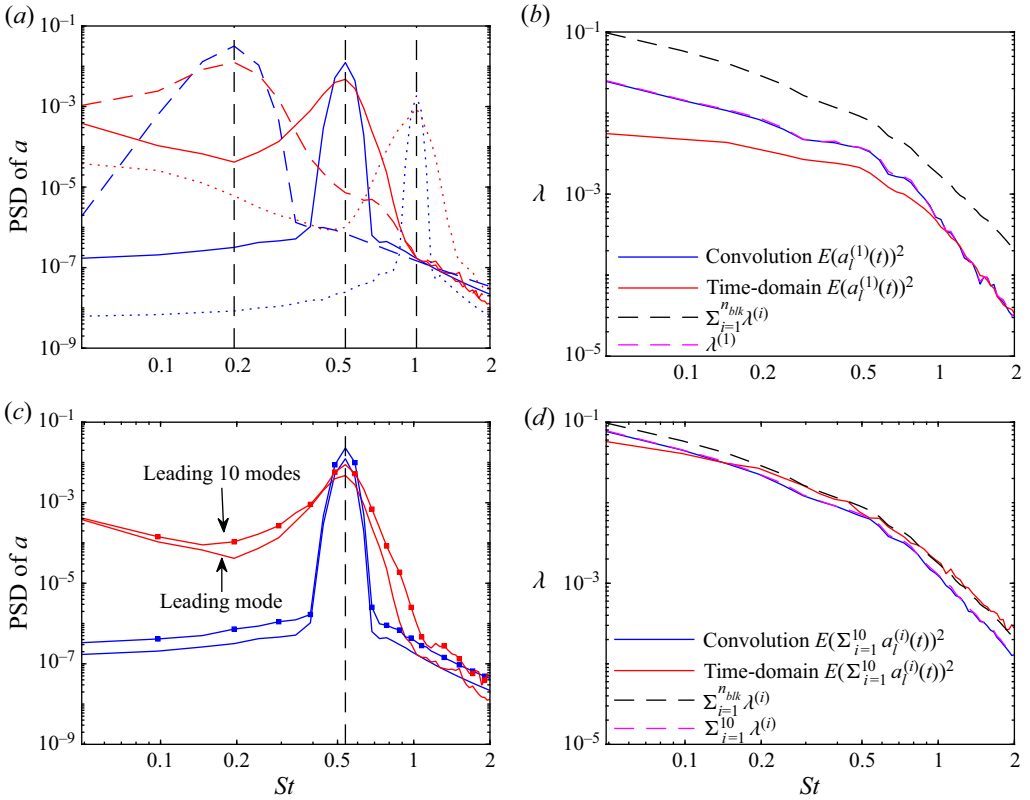


Figure 16. Spectral analysis of the expansion coefficients for the time-domain and convolution approaches: (a) PSD of the individual expansion coefficients of the leading SPOD mode at $St = 0.2, 0.5$ and 1.0 ; (b) expected values of the expansion coefficients of the leading mode; (c) PSD of the expansion coefficients of the leading mode and sum of the first ten leading modes (line with squares) at $St = 0.5$; and (d) expected value of the expansion coefficients for the sum of the first ten modes. In (b,d), the leading, sum of leading 10, and sum of all eigenvalues are shown for comparison.

After establishing the correspondence between the time-domain and convolution approaches, we now compare the spectral characteristics of the two approaches. The PSD of the expansion coefficients associated with the leading mode at $St = 0.2, 0.5$ and 1.0 , for the time and convolution-domain approaches are shown in figure 16(a). As expected the PSD peak at the frequency of the corresponding mode for both approaches. The expansion coefficients computed in the convolution approach exhibit a much narrower peak than those in the time-domain approach. Note that, we cannot expect a sharp spectral peak even for the convolution approach owing to spectral leakage and the modulation of the wave amplitude as seen in figure 15. Figure 16(c) compares the PSD of the leading mode and the ten leading modes at $St = 0.5$, which underlines the dominance of the leading mode at this particular frequency. The expansion coefficients are presented in terms of the spectral energy content in figure 16(b,d). As the expansion coefficients are uncorrelated, and their expected value is equal to the SPOD modal energy,

$$E\{a_f^{(i)} a_f^{(j)}\} = \lambda_f^{(i)} \delta_{ij}, \tag{3.5}$$




Figure 17. Contributions of individual modes to the reconstructed flow field: (a) instantaneous pressure field; (b) contribution of $\phi^{(1)}$ for $St = 0.5$ and (c) for $St = 0.05$; (d) and (e) comparisons of the real part of the pressure field along the lipline ($r = 0.5$) for $St = 0.5$ and $St = 0.05$, respectively. The instantaneous flow field at $t = 1150$ (see figure 15) is shown. Contours in (a–c) are reported on the same colour axis.

it is expected that the convolution approach accurately approximates the eigenvalue spectrum. The blue and magenta lines are almost coincident in figure 16(b,d), thus confirming this conjecture. For the time-domain approach, the expected value for the sum of first 10 modes approximates the total integral PSD for all but the low frequencies $St \geq 0.2$ in figure 16(d). On considering only the first eigenvalue, the time-domain approach underpredicts the eigenvalue spectrum for $St \leq 1.0$, (figure 16b). Note that, this observation does not contradict the observations made in context of figure 4, as the expansion coefficients are obtained from a full basis here, but from a smaller basis in figure 4(d).

After examining the properties of the expansion coefficients, we focus on the spatial composition of the reconstructed flow field in terms of contributions from individual modes in figure 17. The time instant of high energy, previously marked in figure 15, is chosen as an example. The contribution of the leading SPOD modes at two representative frequencies, $St = 0.5$, and 0.05 , to the original flow field is shown in figure 17. Figure 17(b,c) show the modes weighted by their expansion coefficient, $a^{(1)}\phi^{(1)}$ at the corresponding time instant of these two frequencies. Close to the nozzle exit, the LES flow field clearly exhibits a KH-type instability wave. The leading mode at $St = 0.5$ closely resembles this structure. Similarly, the dominant wave pattern with a large wavelength (≈ 5) observed in the flow field is represented by the mode at $St = 0.05$ in a location downstream of the potential core ($12 \lesssim x \lesssim 20$). The real part of the pressure field along the lipline ($r = 0.5$) is compared with the contributions of the leading modes at $St = 0.5$ and $St = 0.05$ in figures 17(d) and 17(e), respectively. It can be seen that the phases of the pressure field and the mode at $St = 0.5$ are aligned in the region where the mode attains its maximum. A weaker but similar kind of phase alignment is also observed for low frequency, $St = 0.05$, despite the disturbed nature of the wavepacket in the region, $12 \lesssim x \lesssim 20$. For the flow field at the current time instant, this also explains that the contribution of the leading mode at $St = 0.05$ is lower than $St = 0.5$, where the KH wavepacket dominates the flow field. These observations confirm that the maximum in the frequency–time diagrams indicate close resemblance of the instantaneous flow field

with the corresponding SPOD modes. Since the SPOD modes are the most energetic structures it is not surprising that the maxima in the frequency–time diagrams indicate the intervals of high energy. Furthermore, as the leading SPOD modes are spatially coherent and contain the most energy at each frequency, we infer that SPOD-based frequency–time analysis can be used to gauge the intermittency of large-scale coherent structures. Here, for brevity, only the convolution approach is shown, but we note that these conclusions also hold for the time-domain approach.

3.4. Prewhitening

As mentioned in § 1, prewhitening is a filtering operation that results in a flat power spectrum, and is commonly used for trend detection in atmospheric and geophysical applications. The goal, therefore, is to use SPOD to scale the data in order to have the same energy at all frequencies. We propose to achieve this goal by rescaling the expansion coefficients of the reconstruction in the frequency domain. The frequency-domain approach is chosen for the same reasons as for denoising in § 3.2. We leverage the fact that the expansion coefficients are uncorrelated, and that their expected value is equal to SPOD modal energy

$$E\{a_f^{(i)} a_f^{(j)}\} = \lambda_f^{(i)} \delta_{ij}. \tag{3.6}$$

We propose two different scalings,

$$\frac{\sqrt{\sum_f \sum_i \lambda_f^{(i)}}}{\sqrt{\lambda_f^{(i)}}} a_f^{(i)} \quad (\text{scaling 1}) \tag{3.7}$$

and

$$\frac{1}{\sqrt{\sum_i \lambda_f^{(i)}}} a_f^{(i)} \quad (\text{scaling 2}) \tag{3.8}$$

to scale the integral mode energy, that is, the sum of all eigenvalues, to one at each frequency. Both methods achieve this goal, but result in different relative scalings of individual SPOD modes. Figures 18(a) and 18(b) show the effect of the two scalings on the SPOD eigenvalue spectrum. The reconstructed flow fields obtained from the expansion coefficients scaled using (3.7) and (3.8) are shown in figures 18(c) and 18(d), respectively. Note that (3.7) collapses all eigenvalues in figure 18(a) to the same value. Scaling 2, on the contrary, preserves both the mode hierarchy and the relative energy content. In comparison to the original flow field shown in figure 6(a), prewhitening emphasises high-frequency structures in the shear layer, whereas it de-emphasises the highly energetic large-scale structures associated with low frequencies downstream of the potential core. This portrayal of the flow field might appear unfamiliar; we emphasise that the objective of prewhitening is not physical interpretation, but pattern identification. Here, for example, the prewhitened pressure fields bring to light the trapped acoustic modes in the potential core. These modes have only recently been described in detail (Schmidt *et al.* 2017b; Towne *et al.* 2017). Previously, they remained largely unnoticed in the analysis of jet data because of their low energy content. An important difference between SPOD-based prewhitening and classical local, point-wise prewhitening techniques is that the SPOD-based approach preserves the spatial coherence of the flow structures identified by the SPOD modes.

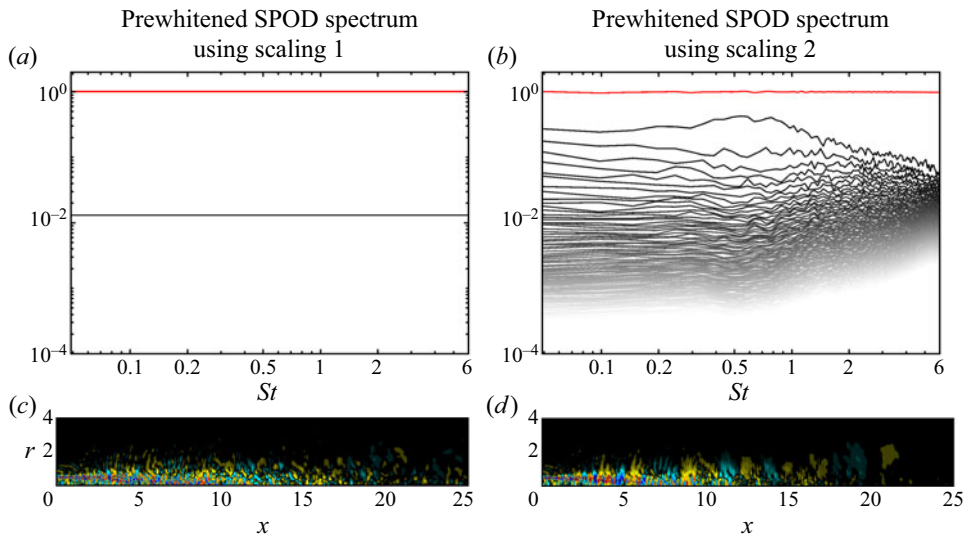


Figure 18. The SPOD-based prewhitening: (a,b) rescaled SPOD spectra; (c,d), instantaneous pressure fields using definitions (3.7) and (3.8), respectively. Red lines indicate the sum of all eigenvalues, that is, the integrated PSD. All lines in (a) collapse owing to the scaling. The original SPOD eigenvalue spectrum can be seen in [figure 2\(a\)](#).

4. Summary and conclusions

Different applications of SPOD including low-rank reconstruction, denoising, prewhitening and frequency–time analysis are demonstrated for the example of LES data of a turbulent jet. A fundamental building block for these applications is the capability to reconstruct the original data from the SPOD. In the frequency domain, this can be accomplished by inverting the SPOD problem (see, e.g. Citriniti & George 2000). We demonstrate that this inversion can be computed either directly in the frequency domain or using a convolution-based strategy. The latter approach becomes a necessity in the context of frequency–time analysis, where the corresponding SPOD problem becomes intractable. As an alternative to frequency-domain reconstruction, we introduce a time-domain approach that is based on the oblique projection of the data onto the SPOD modes. A reduced-order model based on SPOD that uses oblique projection was recently devised by Chu & Schmidt (2020).

Here, we show the complete recovery of the data using all modes and compare with low-dimensional reconstructions. The low-dimensional reconstructions from both approaches accurately capture the integral energy of the segmented data (in the space–time norm). However, the time-varying dynamics (in the purely spatial norm) is only captured by the low-dimensional reconstructions in the time domain. For a fixed number of modes, the time-domain approach captures more of the energy (in both norms). On the downside, the association of the SPOD modes with a single frequency is lost. Instantaneous pressure fields reconstructed in the frequency domain, on the contrary, preserve this monochromatic property of the SPOD modes, but may lack the finer details of the flow-field reconstructions in the time domain. The main advantage of the frequency-domain approach is that it conserves the orthogonality property and the frequency-mode correspondence of the SPOD. The main advantage of the time-domain approach is its optimality in reconstructing the instantaneous flow field with the least possible number of modes.

After establishing the advantages and disadvantages of both approaches, we demonstrate SPOD-based denoising as an application of the frequency-domain approach. As expected, noise is mainly captured by higher SPOD modes at low frequencies and all modes at high frequencies. As a best practice, we propose a hard threshold above the noise floor that is identified from the SPOD eigenvalue spectrum. Significant noise reduction is achieved. At the same time, a substantial amount of energy of the original flow field is retained. Compared to a standard low-pass filter, SPOD-based denoising has the additional advantage of significant storage savings.

Finally, we demonstrate how SPOD-based frequency–time analysis can be used to analyse the intermittency of turbulent flows. Established means of frequency–time analysis such as WT are signal-processing techniques that are applied to one-dimensional time signals. The alternative, SPOD-based approach demonstrated here, provides a global perspective in which spectrograms characterise the temporal evolution of the spatially coherent flow structures represented by the SPOD modes. The SPOD-based frequency–time analysis requires the computation of time-varying expansion coefficients at each time instant, and is computationally intractable in the frequency domain. This problem is mitigated by the convolution-based strategy, which is mathematically equivalent in the limit of the intractable continuously discrete (in time) SPOD problem. This convolution-based approach is compared to the projection-based approach in the time domain. The expansion coefficients calculated from both methods show similar trends. We further demonstrate that a moving average of the spectrogram obtained via oblique projection resembles the spectrogram obtained from the convolution approach. For consistency, the moving time average is directly based on the SPOD windowing. The main advantage of the frequency domain, and therefore the convolution approach, is that it retains the orthogonality property and mode–frequency correspondence of the SPOD. The frequency–time analysis of the jet data confirms the highly intermittent nature of this turbulent flow. In accordance with the SPOD eigenvalue spectrum, it is found that most of the energy is concentrated at low frequencies $St \lesssim 0.2$. A comparison of the total flow energy as a function of time with the spectrograms shows that high-energy events are directly linked to the presence of flow structures resembling the leading SPOD modes. We highlight that this behaviour is not necessarily expected as the SPOD modes represent the most energetic structures only in a statistical sense. From previous work (Schmidt *et al.* 2018), it is well known that SPOD modes often isolate certain, prevailing physical phenomena. The use of SPOD-based frequency–time analysis, hence, provides additional physical insight by indicating time intervals during which a particular mechanism is active.

Based on the results, we recommend the use of the time-domain approach for low-rank reconstruction of individual snapshots, and the frequency-domain approach for denoising and frequency–time analysis. For the latter application the proposed convolution strategy facilitates efficient computation of the time-continuous expansion coefficients. A Matlab code for the convolution-based frequency–time analysis is freely available online.

Acknowledgements. The authors would like to thank the anonymous reviewers for their insightful comments. In particular, we thank the first reviewer for the suggestion regarding time continuity.

Funding. We gratefully acknowledge support from Office of Naval Research grant N00014-20-1-2311.

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

Author ORCIDs.

▣ Akhil Nekkanti <https://orcid.org/0000-0002-2173-8704>;

▣ Oliver T. Schmidt <https://orcid.org/0000-0002-7097-0235>.

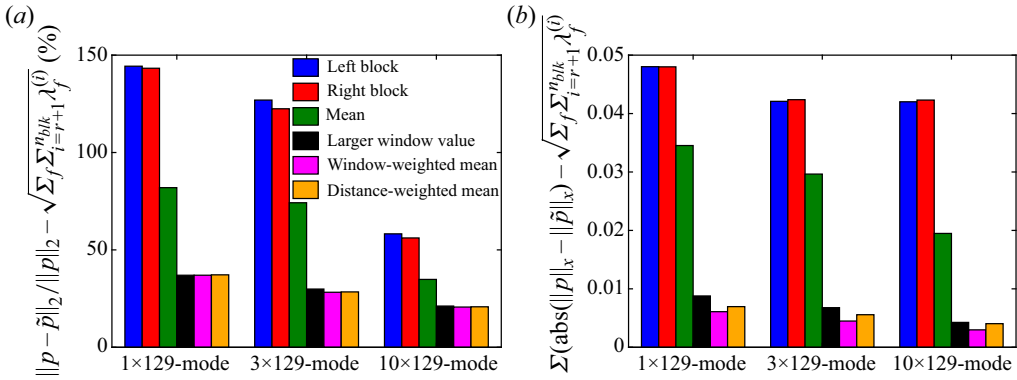


Figure 19. Errors of different ways to compute the frequency-domain reconstruction for 1×129 , 3×129 and 10×129 modes: (a) 2-norm; (b) spatial norm (2.2). Analogous to figure 5, the residual energy of the truncated modes is subtracted for comparison.

Appendix A. Effect of different parameters on the flow-field reconstruction

The segmentation of the data is a crucial step in spectral estimation. Following the original work by Welch (1967), we use an overlap of 50 % between blocks to minimise the variance of the spectral estimate, see § 3 above. The use of overlapping segments, though, results in an ambiguity for the reconstruction, which we may compute using the following:

- (i) the left (previous) block;
- (ii) the right (following) block;
- (iii) the average of the left and right reconstructions;
- (iv) either the left or right reconstruction based on the higher windowing weight;
- (v) the average of the left and right reconstructions weighted by the relative value of Hamming window;
- (vi) the average of the left and right reconstructions weighted by the relative distance to the centres of the overlapping blocks.

Figure 19 compares the errors of low-dimensional reconstructions using 1×129 , 3×129 and 10×129 modes for all the six possibilities. It is found that the reconstruction based on the window-weighted average of the left and right reconstructions produces the smallest error. Based on this finding, this option is used for the frequency-domain reconstruction throughout the paper. Using the window-weighted average approach, the i -th snapshot is reconstructed as

$$q_i \approx \frac{q_j^{(k)} w(j) + q_{j-n_{ovlp}}^{(k+1)} w(|j - n_{ovlp}|)}{w(j) + w(|j - n_{ovlp}|)}, \quad (\text{A1})$$

where $j = i - (k - 1)(n_{fft} - n_{ovlp})$, $i \in [1, n_t]$, $j \in [1, n_{fft}]$ and $k \in [1, n_{blk}]$. From figure 4 (in particular 4b), it becomes apparent that these distinctions only matter for truncated series reconstructions; full-dimensional reconstructions are generally accurate.

The sudden jumps in the local energy of the reconstruction observed in figure 4 are a windowing effect. We demonstrate this by comparison with reconstruction using rectangular windows and no overlap in figure 20. The low-dimensional reconstructions using 1×129 , 3×129 and 10×129 modes; and all modes are shown. Many observations

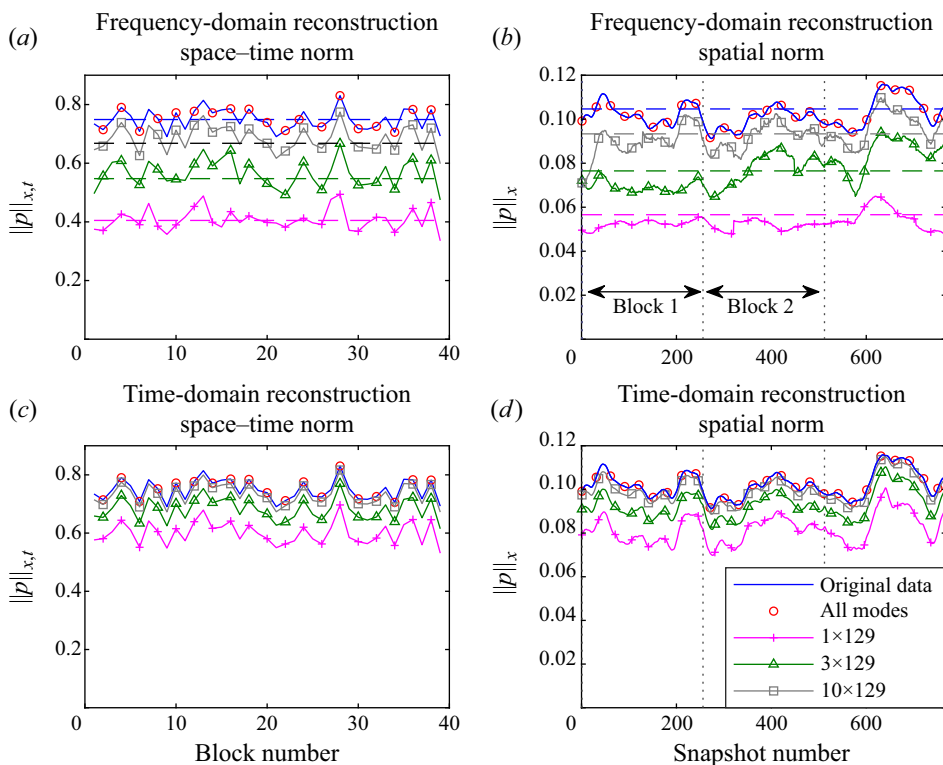


Figure 20. Low-dimensional reconstruction using a rectangular window: frequency-domain reconstruction (a,b), and time-domain reconstruction on the (c,d) in terms of the space-time norm (a,c) and the spatial norm (b,d). The original data (blue lines) is compared with the full reconstructions using all modes, and reconstructions using 10×129 , 3×129 and 1×129 modes. Summed SPOD mode energies are shown as dashed lines. Vertical dotted black lines in (b,d) indicate the non-overlapping blocks. This figure is identical to figure 4, but for a rectangular window and no overlap.

made in the context of figure 4, also hold here: the dimension of the modal bases is directly proportional to its ability to capture the pressure norm of the data and, for a fixed number of modes, the time-domain results in a better approximation of the data than the frequency-domain approach. The most notable difference can be seen between figures 20(b) and 4(b). The windowing effect in the frequency-domain reconstructions is absent if the rectangular window is used. Note in particular the difference during the first few snapshots and near the locations of switching from one block to another (vertical dotted black lines). Despite this advantage in the context of frequency-domain reconstructions for small n_{modes} , rectangular windowing is generally not recommended because of spectral leakage (Schmidt & Colonius 2020).

By analogy with figure 7, we report in figure 21 the SPOD eigenspectra of the 1×129 -mode frequency- and time-domain reconstructions for a rectangular window and $n_{ovlp} = 0$. The SPOD eigenvalue spectra of the full data is also shown for comparison. Only the leading three eigenvalues are shown for clarity. The leading eigenvalue of the frequency-domain and time-domain reconstructions are indistinguishable from the leading eigenvalues of the full data. For the frequency-domain reconstruction, the higher eigenvalue spectra are zero to machine precision, as expected. This indicates that the windowing effect causes the elevation of the higher eigenvalue spectra in figure 7(a). The time-domain reconstruction, on the other hand, is able predict the higher eigenvalues as

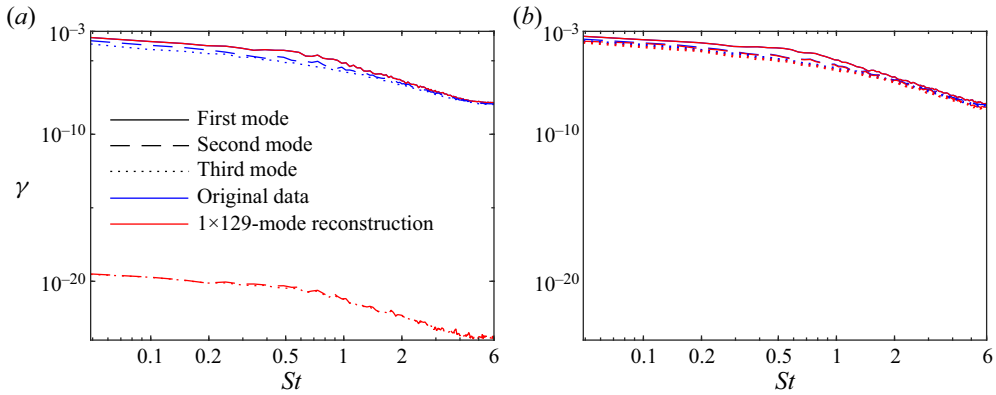


Figure 21. Comparison between SPOD eigenvalue spectra of the original data (blue lines) and 1×129 -mode reconstructions (red lines): (a) frequency domain; (b) time domain. Solid, dashed and dotted lines denote the first, second and third modes, respectively. This figure is identical to figure 7, but for a rectangular window and no overlap.

explained in the context of figure 7(b). This implies that the time-domain reconstruction is much less sensitive to the choice of windowing function.

Appendix B. Projection-based frequency–time analysis: effect of choice of basis and correspondence to convolution-based approach

Oblique projection-based frequency–time analysis is dependent on the choice of modal basis. The two obvious choices of bases for the projection are

- (i) $n_{modes} \times 129$ modes, i.e. only those SPOD modes used in the analysis,

$$\tilde{\Phi} = [\phi_1^{(1)}, \phi_2^{(1)}, \dots, \phi_{n_{fft}}^{(1)}] \tag{B1}$$

or

- (ii) $n_{blk} \times 129$ modes, i.e. all available SPOD modes,

$$\tilde{\Phi} = [\phi_1^{(1)}, \phi_1^{(2)}, \dots, \phi_1^{(n_{blk})}, \phi_2^{(1)}, \phi_2^{(2)}, \dots, \phi_2^{(n_{blk})}, \dots, \phi_{n_{fft}}^{(1)}, \phi_{n_{fft}}^{(2)}, \dots, \phi_{n_{fft}}^{(n_{blk})}]. \tag{B2}$$

Owing to the non-orthogonality of these modes in the spatial norm, these two choices will result in different outcomes. Shown in figure 22(a) is the frequency–time diagram for $n_{modes} = 1$, that is a 1×129 -mode basis containing only the leading mode at each frequency. A fundamentally different behaviour from that in figure 12 is observed. The diagram exhibits a banded structure, and, in contrast to the reference diagram based on all SPOD modes, the majority of maxima is not found in the low-frequency regime, $St \lesssim 0.2$. To understand this difference, the PSD of the expansion coefficient associated with the leading mode at $St = 0.5$ is shown in figure 22(b). The expansion coefficient computed with the 1×129 -modal basis exhibits a much broader peak than the one computed with full basis. This behaviour indicates a loss of the mode–frequency correspondence for the heavily truncated basis. Next, both approaches are compared by taking the convolution-based expansion coefficient as the reference. The cross-correlation of the expansion coefficients from both approaches with the reference signal from the convolution approach are shown in figure 22(c). The expansion coefficient computed




Figure 22. The SPOD-based frequency–time analysis obtained using the time-domain approach with 1×129 -modal basis: (a) frequency–time diagram of the leading mode at each frequency; (b) PSD of the individual expansion coefficients of the leading SPOD mode at $St = 0.5$ for the time-domain approach with the 1×129 basis (blue line) and $n_{blk} \times 129$ basis (red line); (c) cross-correlation of the expansion coefficient in the time-domain approach using the 1×129 basis (blue line) and the $n_{blk} \times 129$ basis (red dotted line) with the convolution approach.




Figure 23. The SPOD-based frequency–time diagrams obtained using the moving mean of the time-domain oblique projection based approach: (a) first 10 modes at each frequency; (b) leading mode at each frequency. The moving mean uses as weights the same Hamming windowing function as the SPOD.

using the 1×129 -modal basis exhibits a much lower correlation with the reference. We conclude from this analysis that the time-domain approach should be conducted using all SPOD modes, as it yields a more accurate description of the intermittency of the coherent structure represented by the SPOD modes.

In figure 15, we demonstrated that the expansion coefficients computed from a moving average of the time-domain approach resemble those from the convolution approach. For further evidence, we show in figure 23 the frequency–time diagrams obtained by taking the moving mean, at each frequency, of the time-domain diagram previously shown in figure 12. The outcome should be compared to the frequency–time diagrams obtained using the convolution approach, i.e. figure 13. It is observed that the frequency–time

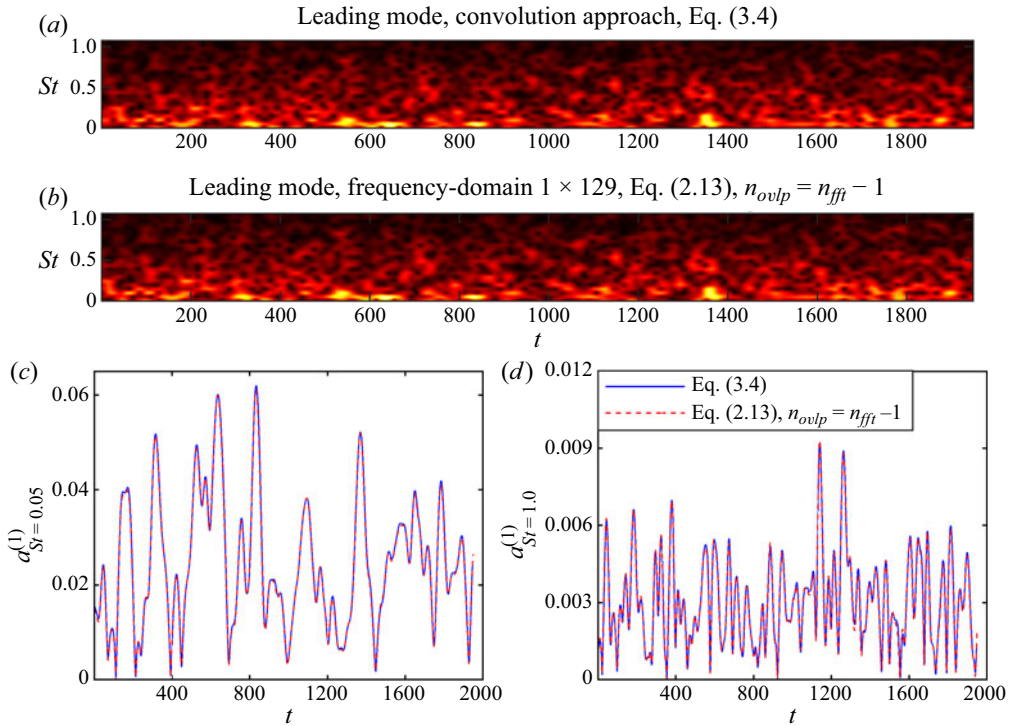


Figure 24. Comparison between two strategies for SPOD-based frequency–time analysis: frequency–time diagrams for (a) the convolution approach (SPOD modes precomputed with 50%, overlap) and (b) the frequency-domain approach using $n_{ovlp} = n_{fft} - 1$; expansion coefficients for (c) $St = 0.05$ and (d) $St = 1.0$.

diagrams are qualitatively very similar. The effect of taking the moving average is mainly visible at higher frequencies, where it leads to minor loss of detail. For qualitative flow analysis, we therefore conclude that the moving average of the time-domain approach can well be used to approximate the much more computationally involved convolution approach.

Appendix C. Frequency–time analysis based on frequency-domain approach: effect of overlap and correspondence to convolution-based approach

Figure 24 demonstrates the similarity between the frequency–time diagrams obtained using the convolution method and the frequency-domain approach. For the convolution method, (3.4), a basis of precomputed SPOD modes with 50% overlap was used. To obtain time-continuous expansion coefficients using the frequency-domain approach, we require $n_{ovlp} = n_{fft} - 1$. As this is computationally intractable, the data (only here) were reduced to every third grid point in the streamwise and radial directions. The time traces of the individual expansion coefficients for $St = 0.05$ and 1.0 are reported in figure 24(c,d).

REFERENCES

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. *Theor. Comput. Fluid Dyn.* **2** (5–6), 339–352.

- BALE, S.D., KELLOGG, P.J., MOZER, F.S., HORBURY, T.S. & REME, H. 2005 Measurement of the electric fluctuation spectrum of magnetohydrodynamic turbulence. *Phys. Rev. Lett.* **94** (21), 215002.
- BOASHASH, B. 1988 Note on the use of the wigner distribution for time-frequency signal analysis. *IEEE Trans. Acoust. Speech Signal Process.* **36** (9), 1518–1521.
- BRÈS, G.A., HAM, F.E., NICHOLS, J.W. & LELE, S.K. 2017 Unstructured large-eddy simulations of supersonic jets. *AIAA J.* **55** (4), 1164–1184.
- BRÈS, G.A., JORDAN, P., JAUNET, V., LE RALLIC, M., CAVALIERI, A.V.G., TOWNE, A., LELE, S.K., COLONIUS, T. & SCHMIDT, O.T. 2018 Importance of the nozzle-exit boundary-layer state in subsonic turbulent jets. *J. Fluid Mech.* **851**, 83–124.
- BRÈS, G.A. & LELE, S.K. 2019 Modelling of jet noise: a perspective from large-eddy simulations. *Phil. Trans. R. Soc. A* **377** (2159), 20190081.
- BRINDISE, M.C., CHIASTRA, C., BURZOTTA, F., MIGLIAVACCA, F. & VLACHOS, P.P. 2017 Hemodynamics of stent implantation procedures in coronary bifurcations: an in vitro study. *Ann. Biomed. Engng* **45** (3), 542–553.
- BRINDISE, M.C. & VLACHOS, P.P. 2017 Proper orthogonal decomposition truncation method for data denoising and order reduction. *Exp. Fluids* **58** (4), 28.
- BROWN, E.D., BUCHSBAUM, S.B., HALL, R.E., PENHUNE, J.P., SCHMITT, K.F., WATSON, K.M. & WYATT, D.C. 1989 Observations of a nonlinear solitary wave packet in the Kelvin wake of a ship. *J. Fluid Mech.* **204**, 263–293.
- CAMUSSI, R. 2002 Coherent structure identification from wavelet analysis of particle image velocimetry data. *Exp. Fluids* **32** (1), 76–86.
- CAMUSSI, R. & GUI, G. 1997 Orthonormal wavelet decomposition of turbulent flows: intermittency and coherent structures. *J. Fluid Mech.* **348**, 177–199.
- CHARONKO, J., KARRI, S., SCHMIEG, J., PRABHU, S. & VLACHOS, P. 2010 In vitro comparison of the effect of stent configuration on wall shear stress using time-resolved particle image velocimetry. *Ann. Biomed. Engng* **38** (3), 889–902.
- CHU, T. & SCHMIDT, O.T. 2020 Stochastic SPOD-Galerkin model of a turbulent jet. *Theor. Comput. Fluid Dyn.* (submitted) [arXiv:2012.02902](https://arxiv.org/abs/2012.02902).
- 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.
- COHEN, L. 1995 *Time-Frequency Analysis*, vol. 778. Prentice Hall.
- DELVILLE, J. 1994 Characterization of the organization in shear layers via the proper orthogonal decomposition. *Appl. Sci. Res.* **53** (3), 263–281.
- DISCETTI, S., NATALE, A. & ASTARITA, T. 2013 Spatial filtering improved tomographic PIV. *Exp. Fluids* **54** (4), 1505.
- FARGE, M. 1992 Wavelet transforms and their applications to turbulence. *Annu. Rev. Fluid Mech.* **24** (1), 395–458.
- FARGE, M., PELLEGRINO, G. & SCHNEIDER, K. 2001 Coherent vortex extraction in 3D turbulent flows using orthogonal wavelets. *Phys. Rev. Lett.* **87** (5), 054501.
- FARGE, M., SCHNEIDER, K. & KEVLAHAN, N. 1999 Non-Gaussianity and coherent vortex simulation for two-dimensional turbulence using an adaptive orthogonal wavelet basis. *Phys. Fluids* **11** (8), 2187–2201.
- FORE, L.B., TUNG, A.T., BUCHANAN, J.R. & WELCH, J.W. 2005 Nonlinear temporal filtering of time-resolved digital particle image velocimetry data. *Exp. Fluids* **39** (1), 22–31.
- FREUND, J.B. & COLONIUS, T. 2009 Turbulence and sound-field POD analysis of a turbulent jet. *Intl J. Aeroacoust.* **8** (4), 337–354.
- GAMARD, S., GEORGE, W.K., JUNG, D. & WOODWARD, S. 2002 Application of a “slice” proper orthogonal decomposition to the far field of an axisymmetric turbulent jet. *Phys. Fluids* **14** (7), 2515–2522.
- GAMARD, S., JUNG, D. & GEORGE, W.K. 2004 Downstream evolution of the most energetic modes in a turbulent axisymmetric jet at high Reynolds number. Part 2. The far-field region. *J. Fluid Mech.* **514**, 205–230.
- GHATE, A.S., TOWNE, A. & LELE, S.K. 2020 Broadband reconstruction of inhomogeneous turbulence using spectral proper orthogonal decomposition and Gabor modes. *J. Fluid Mech.* **888**, R1.
- GLAUSER, M.N. & GEORGE, W.K. 1992 Application of multipoint measurements for flow characterization. *Expl Therm. Fluid Sci.* **5** (5), 617–632.
- GLAUSER, M.N., LEIB, S.J. & GEORGE, W.K. 1987 Coherent structures in the axisymmetric turbulent jet mixing layer. In *Turbulent Shear Flows 5*, pp. 134–145. Springer.
- 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.

- GORDEYEV, S.V. & THOMAS, F.O. 2002 Coherent structure in the turbulent planar jet. Part 2. Structural topology via POD eigenmode projection. *J. Fluid Mech.* **460**, 349–380.
- GU, D. & PHILANDER, S.G.H. 1995 Secular changes of annual and interannual variability in the tropics during the past century. *J. Clim.* **8** (4), 864–876.
- GUDMUNDSSON, K. & COLONIUS, T. 2011 Instability wave models for the near-field fluctuations of turbulent jets. *J. Fluid Mech.* **689**, 97–128.
- HARRIS, F.J. 1978 On the use of windows for harmonic analysis with the discrete Fourier transform. *Proc. IEEE* **66** (1), 51–83.
- HELLSTRÖM, L.H.O., GANAPATHISUBRAMANI, B. & SMITS, A.J. 2015 The evolution of large-scale motions in turbulent pipe flow. *J. Fluid Mech.* **779**, 701–715.
- HELLSTRÖM, L.H.O., MARUSIC, I. & SMITS, A.J. 2016 Self-similarity of the large-scale motions in turbulent pipe flow. *J. Fluid Mech.* **792**.
- HELLSTRÖM, L.H.O. & SMITS, A.J. 2014 The energetic motions in turbulent pipe flow. *Phys. Fluids* **26** (12), 125102.
- HELLSTRÖM, L.H.O. & SMITS, A.J. 2017 Structure identification in pipe flow using proper orthogonal decomposition. *Phil. Trans. R. Soc. A* **375** (2089), 20160086.
- HOLMES, P., LUMLEY, J.L., BERKOOZ, G. & ROWLEY, C.W. 2012 *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge University Press.
- HUANG, N.E., SHEN, Z., LONG, S.R., WU, M.C., SHIH, H.H., ZHENG, Q., YEN, N.C., TUNG, C.C. & LIU, H.H. 1998 The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis. *Proc. R. Soc. Lond. A* **454** (1971), 903–995.
- IQBAL, M.O. & THOMAS, F.O. 2007 Coherent structure in a turbulent jet via a vector implementation of the proper orthogonal decomposition. *J. Fluid Mech.* **571**, 281–326.
- IZATT, J.A., KULKARNI, M.D., YAZDANFAR, S., BARTON, J.K. & WELCH, A.J. 1997 In vivo bidirectional color doppler flow imaging of picoliter blood volumes using optical coherence tomography. *Opt. Lett.* **22** (18), 1439–1441.
- JOHANSSON, P.B.V. & GEORGE, W.K. 2006a The far downstream evolution of the high-Reynolds-number axisymmetric wake behind a disk. Part 1. Single-point statistics. *J. Fluid Mech.* **555**, 363–385.
- JOHANSSON, P.B.V. & GEORGE, W.K. 2006b The far downstream evolution of the high-Reynolds-number axisymmetric wake behind a disk. Part 2. Slice proper orthogonal decomposition. *J. Fluid Mech.* **555**, 387–408.
- JOHANSSON, P.B.V., GEORGE, W.K. & WOODWARD, S.H. 2002 Proper orthogonal decomposition of an axisymmetric turbulent wake behind a disk. *Phys. Fluids* **14** (7), 2508–2514.
- JUNG, D., GAMARD, S. & GEORGE, W.K. 2004 Downstream evolution of the most energetic modes in a turbulent axisymmetric jet at high Reynolds number. Part 1. The near-field region. *J. Fluid Mech.* **514**, 173–204.
- KHALIQ, M.N., OUARDA, T.B.M.J., GACHON, P., SUSHAMA, L. & ST-HILAIRE, A. 2009 Identification of hydrological trends in the presence of serial and cross correlations: A review of selected methods and their application to annual flow regimes of Canadian rivers. *J. Hydrol.* **368** (1–4), 117–130.
- KOSTAS, J., SORIA, J. & CHONG, M.S. 2005 A comparison between snapshot pod analysis of PIV velocity and vorticity data. *Exp. Fluids* **38** (2), 146–160.
- LACOMBE, G., MCCARTNEY, M. & FORKUOR, G. 2012 Drying climate in Ghana over the period 1960–2005: evidence from the resampling-based Mann-Kendall test at local and regional levels. *Hydrol. Sci. J.* **57** (8), 1594–1609.
- LUMLEY, J.L. 1967 The structure of inhomogeneous turbulent flows. In *Atmospheric Turbulence and Radio Wave Propagation* (ed. A.M. Yaglom & V.I. Tatarski), pp. 166–178. Nauka.
- LUMLEY, J.L. 1970 *Stochastic Tools in Turbulence*. Courier Corporation.
- MENEVEAU, C., LUND, T.S. & MOIN, P. 1992 Search for subgrid scale parameterization by projection pursuit regression. In *Proceedings of Summer Program*, pp. 61–81. Stanford University.
- MEYERS, S.D., KELLY, B.G. & O'BRIEN, J.J. 1993 An introduction to wavelet analysis in oceanography and meteorology: with application to the dispersion of Yanai waves. *Mon. Weath. Rev.* **121** (10), 2858–2866.
- MURALIDHAR, S.D., PODVIN, B., MATHELIN, L. & FRAIGNEAU, Y. 2019 Spatio-temporal proper orthogonal decomposition of turbulent channel flow. *J. Fluid Mech.* **864**, 614–639.
- NEKKANTI, A. & SCHMIDT, O.T. 2021 Modal analysis of acoustic directivity in turbulent jets. *AIAA J.* **59** (1), 1–12.
- NIDHAN, S., CHONGSIRIPINYO, K., SCHMIDT, O.T. & SARKAR, S. 2020 Spectral proper orthogonal decomposition analysis of the turbulent wake of a disk at $Re = 50\,000$. *Phys. Rev. Fluids* **5** (12), 124606.
- ONORATO, M., CAMUSSI, R. & IUSO, G. 2000 Small scale intermittency and bursting in a turbulent channel flow. *Phys. Rev. E* **61** (2), 1447.

- PICARD, C. & DELVILLE, J. 2000 Pressure velocity coupling in a subsonic round jet. *Intl J. Heat Fluid Flow* **21** (3), 359–364.
- PICKERING, E., RIGAS, G., NOGUEIRA, P.A.S., CAVALIERI, A.V.G., SCHMIDT, O.T. & COLONIUS, T. 2019 Lift-up, Kelvin-Helmholtz and Orr mechanisms in turbulent jets. *J. Fluid Mech.* **896**, A2.
- RAIOLA, M., DISCETTI, S. & IANIRO, A. 2015 On PIV random error minimization with optimal POD-based low-order reconstruction. *Exp. Fluids* **56** (4), 75.
- RODIONOV, S.N. 2006 Use of prewhitening in climate regime shift detection. *Geophys. Res. Lett.* **33** (12).
- 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., *et al.* 2009 Spectral analysis of nonlinear flows. *J. Fluid Mech.* **641** (1), 115–127.
- SAMIMY, M., DEBIASI, M., CARABALLO, E., SERRANI, A., YUAN, X., LITTLE, J. & MYATT, J.H. 2007 Feedback control of subsonic cavity flows using reduced-order models. *J. Fluid Mech.* **579**, 315.
- SCHMID, P.J. 2010 Dynamic mode decomposition of numerical and experimental data. *J. Fluid Mech.* **656**, 5–28.
- SCHMIDT, O.T. & COLONIUS, T. 2020 Guide to spectral proper orthogonal decomposition. *AIAA J.* **58** (3), 1–11.
- SCHMIDT, O.T., COLONIUS, T. & BRÉS, G.A. 2017a Wavepacket intermittency and its role in turbulent jet noise. *AIAA Paper* 2017-0686.
- SCHMIDT, O.T., TOWNE, A., COLONIUS, T., CAVALIERI, A.V.G., JORDAN, P. & BRÉS, G.A. 2017b 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. & BRÉS, G.A. 2018 Spectral analysis of jet turbulence. *J. Fluid Mech.* **855**, 953–982.
- SCIACCHITANO, A. & SCARANO, F. 2014 Elimination of PIV light reflections via a temporal high pass filter. *Meas. Sci. Technol.* **25** (8), 084009.
- SERINALDI, F. & KILSBY, C.G. 2015 Stationarity is undead: uncertainty dominates the distribution of extremes. *Adv. Water. Resour.* **77**, 17–36.
- SIROVICH, L. 1987 Turbulence and the dynamics of coherent structures. I. Coherent structures. *Q. Appl. Maths* **45** (3), 561–571.
- SON, S.Y. & KIHM, K.D. 2001 Evaluation of transient turbulent flow fields using digital cinematographic particle image velocimetry. *Exp. Fluids* **30** (5), 537–550.
- STEWART, K.C. & VLACHOS, P.P. 2012 Vortex rings in radially confined domains. *Exp. Fluids* **53** (4), 1033–1044.
- STOCKWELL, R.G., MANSINHA, L. & LOWE, R.P. 1996 Localization of the complex spectrum: the S transform. *IEEE Trans. Signal Process.* **44** (4), 998–1001.
- 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.T.M., ROWLEY, C.W., COLONIUS, T., MCKEON, B.J., SCHMIDT, O.T., GORDEYEV, S., THEOFILIS, V. & UKEILEY, L.S. 2017 Modal analysis of fluid flows: an overview. *AIAA J.* **55** (12), 4013–4041.
- TINNEY, C.E., GLAUSER, M.N. & UKEILEY, L.S. 2008a Low-dimensional characteristics of a transonic jet. Part 1. Proper orthogonal decomposition. *J. Fluid Mech.* **612**, 107–141.
- TINNEY, C.E., UKEILEY, L.S. & GLAUSER, M.N. 2008b Low-dimensional characteristics of a transonic jet. Part 2. Estimate and far-field prediction. *J. Fluid Mech.* **615**, 53–92.
- TISSOT, G., LAJÚS, F.C. JR., Cavalieri, A.V.G. & Jordan, P. 2017 Wave packets and Orr mechanism in turbulent jets. *Phys. Rev. Fluids* **2** (9), 093901.
- TOWNE, A., CAVALIERI, A.V.G., JORDAN, P., COLONIUS, T., SCHMIDT, O., JAUNET, V. & BRÉS, G.A. 2017 Acoustic resonance in the potential core of subsonic jets. *J. Fluid Mech.* **825**, 1113–1152.
- TOWNE, A. & LIU, P. 2019 Time-frequency analysis of intermittent coherent structures in turbulent flows. *Bull. Am. Phys. Soc.* **64**.
- TOWNE, A., SCHMIDT, O.T. & COLONIUS, T. 2018 Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis. *J. Fluid Mech.* **847**, 821–867.
- 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), 1118–1134.
- VÉTEL, J., GARON, A. & PELLETIER, D. 2011 Denoising methods for time-resolved PIV measurements. *Exp. Fluids* **51** (4), 893–916.
- VON STORCH, H. 1999 Misuses of statistical analysis in climate research. In *Analysis of climate variability*, pp. 11–26. Springer.

SPOD for analysis of turbulent flows

- 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.
- ZHANG, X., HARVEY, K.D., HOGG, W.D. & YUZYK, T.R. 2001 Trends in Canadian streamflow. *Water Resour. Res.* **37** (4), 987–998.
- ZHANG, Y., GUO, Z., WANG, W., HE, S., LEE, T. & LOEW, M. 2003 A comparison of the wavelet and short-time Fourier transforms for Doppler spectral analysis. *Med. Engng Phys.* **25** (7), 547–557.